44835

National Library
of Canada

Bibliothèque nationale
du Canada

Canadian Theses Division

Division des thèses canadiennes

Ottawa, Canada
K1A 0N4

PERMISSION TO MICROFILM — AUTORISATION DE MICROFILMER

- Please print or type — Écrire en lettres moulées ou dactylographier

Full Name of Author — Nom complet de l'auteur

VERA SYLVIA TOWNSEND

Date of Birth — Date de naissance MAY 27, 1954	Country of Birth — Lieu de naissance CANADA

Permanent Address — Résidence fixe

#31, 10824 - 85 AVENUE
EDMONTON, ALBERTA T6E 2L1

Title of Thesis — Titre de la thèse

PERFORMANCE MEASUREMENT OF COMPUTER COMMUNICATION NETWORKS

University — Université

UNIVERSITY OF ALBERTA

Degree for which thesis was presented — Grade pour lequel cette thèse fut présentée

MASTER OF SCIENCE

Year this degree conferred — Année d'obtention de ce grade 1980	Name of Supervisor — Nom du directeur de thèse PROF. URSULA MAYDELL

Permission is hereby granted to the NATIONAL LIBRARY OF CANADA to microfilm this thesis and to lend or sell copies of the film.

The author reserves other publication rights, and neither the thesis nor extensive extracts from it may be printed or otherwise reproduced without the author's written permission.

L'autorisation est, par la présente, accordée à la BIBLIOTHÈQUE NATIONALE DU CANADA de microfilmer cette thèse et de prêter ou de vendre des exemplaires du film.

L'auteur se réserve les autres droits de publication; ni la thèse ni de longs extraits de celle-ci ne doivent être imprimés ou autrement reproduits sans l'autorisation écrite de l'auteur.

Date APRIL 25, 1980	Signature Vera Townsend

NL-91 (4/77)

National Library of Canada
Collections Development Branch

Canadian Theses on
Microfiche Service

Bibliothèque nationale du Canada
Direction du développement des collections

Service des thèses canadiennes
sur microfiche

NOTICE

The quality of this microfiche is heavily dependent upon the quality of the original thesis submitted for microfilming. Every effort has been made to ensure the highest quality of reproduction possible.

If pages are missing, contact the university which granted the degree.

Some pages may have indistinct print especially if the original pages were typed with a poor typewriter ribbon or if the university sent us a poor photocopy.

Previously copyrighted materials (journal articles, published tests, etc.) are not filmed.

Reproduction in full or in part of this film is governed by the Canadian Copyright Act, R.S.C. 1970, c. C-30. Please read the authorization forms which accompany this thesis.

THIS DISSERTATION
HAS BEEN MICROFILMED
EXACTLY AS RECEIVED

AVIS

La qualité de cette microfiche dépend grandement de la qualité de la thèse soumise au microfilmage. Nous avons tout fait pour assurer une qualité supérieure de reproduction.

S'il manque des pages, veuillez communiquer avec l'université qui a conféré le grade.

La qualité d'impression de certaines pages peut laisser à désirer, surtout si les pages originales ont été dactylographiées à l'aide d'un ruban usé ou si l'université nous a fait parvenir une photocopie de mauvaise qualité.

Les documents qui font déjà l'objet d'un droit d'auteur (articles de revue, examens publiés, etc.) ne sont pas microfilmés.

La reproduction, même partielle, de ce microfilm est soumise à la Loi canadienne sur le droit d'auteur, SRC 1970, c. C-30. Veuillez prendre connaissance des formules d'autorisation qui accompagnent cette thèse.

LA THÈSE A ÉTÉ
MICROFILMÉE TELLE QUE
NOUS L'AVONS REÇUE

Ottawa, Canada
K1A 0N4

NL-339 (Rev. 8/80)

THE UNIVERSITY OF ALBERTA

Performance Measurement of Computer Communication Networks

by

© Vera Sylvia Townsend

A THESIS
SUBMITTED TO THE FACULTY OF GRADUATE STUDIES AND RESEARCH
IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE DEGREE
OF Master of Science

DEPARTMENT OF Computing Science

EDMONTON, ALBERTA
Spring, 1980

THE UNIVERSITY OF ALBERTA

FACULTY OF GRADUATE STUDIES AND RESEARCH

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research, for acceptance, a thesis entitled Performance Measurement of Computer Communication Networks submitted by Vera Sylvia Townsend in partial fulfilment of the requirements for the degree of Master of Science.

....

Supervisor

L. N. Jackson

Chin

McLeod

Date April 22, 1980

Abstract

The desire to enhance access to and foster sharing of data processing resources has stimulated the development of computer networks. Performance studies evaluate the efficiency of networks for the purpose of system improvement and expansion.

The role of measurement in the evaluation of computer networks is the central theme of this thesis. Results and techniques cited in the recent literature are examined, compared and contrasted.

A computerized bibliographic database, compiled by the author, on network design and performance analysis is discussed. An overview of different network implementations and international protocol standardization activities establishes the terminology framework for the study.

The goals and objectives of a measurement experiment and the requirements of a measurement facility are identified. The hardware and software tools from various networks are compared.

A model for the characterization of the user input process is developed, and observed values of the workload variables are noted.

Throughput, delay, fairness, stability, protocol overhead, reliability and availability are the network performance measures assessed. Observations from different networks using real or artificially generated traffic are summarized, and presented in standard units. Conclusions

about the relative performance of the various networks are then drawn.

Finally, the role of performance measurement and the results obtained are interpreted in terms of computer communication in the future.

Acknowledgment

I wish to thank Professor Ursula Maydell, my supervisor, for her advice, criticism and guidance during the research and preparation of this thesis. In addition, I wish to acknowledge the support and encouragement given me by my family, relatives, friends, and colleagues at the Alberta Research Council.

Table of Contents

Chapter Page

1. Introduction 1

2. Computer Network Design and Performance Database 6

  2.1 Guidelines for Document Selection 7

  2.2 SPIRES Database Management System 8

  2.3 The NETWORKS Database 9

  2.4 Summary 13

3. Definitions and the Role of Standardization 14

  3.1 Definitions 14

  3.2 Types of Networks 16

  3.3 Open System Interconnection 18

  3.4 Summary 26

4. Network Protocols and Services 27

  4.1 Packet Switched Networks 27

    4.1.1 ARPANET Network 30

    4.1.2 CYCLADES Network 33

    4.1.3 DATAPAC Network 34

    4.1.4 IBM Experimental Satellite Link Network 35

    4.1.5 KUIPNET Network 35

  4.2 Multiaccess Networks 36

    4.2.1 ALOHANET Network 38

    4.2.2 ETHERNET Network 39

    4.2.3 Packet Radio Experimental (PRNET) Network 40

    4.2.4 SATNET Network 41

  4.3 Summary 46

5. Network Measurement Planning 47

5.1 Network Measurement Goals and Objectives 47

5.2 Requirements of a Measurement Facility 53

5.3 Measurement Data Types 53

5.4 Summary 55

6. Network Measurement Tools 56

6.1 Artificial Traffic Generators 57

6.2 Software Monitors 62

6.2.1 Cumulative Statistics 63

6.2.2 Snapshot Statistics 68

6.2.3 Trace Statistics 68

6.3 Hardware Monitors 71

6.4 Network Control 72

6.4.1 Experimental Control 73

6.4.2 Network Monitoring 74

6.5 Summary, Conclusions and Open Problems 76

7. Network Input Process Characterization 78

7.1 Data Stream Model 78

7.2 Traffic Measurement Experiments 84

7.2.1 Message and Packet Size Distribution 84

7.2.2 Message Interarrival Time Distribution ...86

7.2.3 Path Length Distribution 88

7.2.4 Traffic Distribution 89

7.3 Summary, Conclusions and Open Problems 92

8. Network Throughput and Delay 95

8.1 Packet Switched Network Performance 97

8.2 Multiaccess Network Performance 102

8.3 Internetwork Performance 107

8.4 Summary, Conclusions and Open Problems 108

9. Fairness, Stability and Protocol Efficiency 111

9.1 Fairness 111

9.2 System Stability 115

9.3 Protocol Efficiency 116

9.4 Summary, Conclusions and Open Problems 121

10. Network Reliability 124

10.1 Sensitivity to Channel Noise 125

10.2 Availability 129

10.3 Summary, Conclusions and Open Problems 131

11. Summary, Conclusions and Future Research 133

11.1 Standardization 135

11.2 New Network Services 136

11.3 Technological Advances 137

11.4 Network Measurement Tools 139

11.5 Network User Workload 139

11.6 Network Performance 140

11.7 Epilog 142

Annotated Bibliography 143

Appendix A -- Overview of SPIRES 169
A.1 SPIRES Command and Definition Languages 170
A.2 SPIRES Program Modules 172
A.3 Physical Organization 173
A.4 Privacy and Security 175
A.5 Database Recovery 177
A.6 File Management Tools 177
A.7 SPIRES Documentation 178

Appendix B -- Database Design Documentation 180
B.1 File Definition 180
B.2 Format Definitions 192
B.3 Protocol Definitions 195

List of Tables

Table Page

1. International Communication Standards and Activities 24

2. Non X.25 Public Packet Switched Networks 28

3. X.25 Compatible Public Packet Switched Networks 31

4. Multiaccess Networks 37

5. Characterization of User Traffic Types 59

6. Artificial Traffic Generator Parameters 61

7. Software Monitor Statistics and Data Measurement Types 64

8. Cumulative Statistics Collected by Software Monitors 66

9. Snapshot Statistics Collected by Software Monitors 69

10. Trace Statistics Collected by Software Monitors 70

11. Statistics Collected for Network Monitoring 75

12. Message and Packet Size Measurements 85

13. Message Interarrival Time Measurements 87

14. Traffic Distribution Measurements 90

15. Throughput Measurements in Packet Switched Networks 98

List of Tables (cont'd)

Table — Page

16. Delay Measurements in Packet Switched Networks 99

17. Throughput Measurements in Multiaccess Networks 103

18. Delay Measurements in Multiaccess Networks 104

19. Fairness Measurements 113

20. Stability Measurements 117

21. Protocol Overhead Measurements 119

22. Reliability Measurements 126

23. Availability Measurements 130

List of Figures

Figure Page

1. Process to Process Communication Within and Between Workstations in Open System Interconnection 21

2. Accepted International Standards in Relation to Open System Interconnection 23

3. User to Computer Data Stream Model 81

4. Simplified SPIRES B-tree Structure 174

1. Introduction

The tremendous power of data processing systems today is witnessed by the diversity of applications supported by specialized hardware and software on micro-, mini- and large time sharing computer systems. The efforts to enhance access to and foster sharing of these unique resources heralded a new era in computer technology through the development of communication networks.

The objective of this thesis is to examine, compare and contrast measurements and measurement techniques used in the performance evaluation of experimental or operational computer communication networks reported in the recent literature.

Computer networks today range in structure and sophistication from specialized networks designed for specialized tasks, to general purpose networks designed to provide a uniform mechanism for connecting together heterogeneous host computer systems. In all cases, however, the network consists of a collection of switching computers called nodes, and a communication medium.

The design of a network centers around two fundamental approaches to allocation of communication channel bandwidth: preallocation, and dynamic allocation.

In preallocation systems, such as private leased line services, a fixed amount of bandwidth is assigned to each user for the duration of the session. This technique is also known as circuit switching.

Dynamic allocation systems, on the other hand, do not schedule bandwidth over the entire message path from source to destination, but instead assign channel resources needed as the message travels from link to link in the network. Usually the user data is divided into units known as packets prior to transmission over the communication channel: this technique is referred to as packet switching. Processing power and buffer storage at each switch controls and directs the flow of packets.

The economic tradeoff [ROBE78] of packet switching versus circuit switching for computer communication is simple: if the communication bandwidth is cheap, use circuit switching; if computer processing power is cheap, use packet switching. With rapidly declining costs of hardware, and corresponding increase in processing power, packet switching today is the most cost effective solution to linking users and computers together.

A cost effective network, however, is not necessarily an efficient network. It is the role of performance studies to evaluate the behavioral characteristics and optimize utilization of system components during the design, implementation and operational phases.

Beginning with network design, system modelling and simulation studies work hand in hand to synthesize the design variables [KLEI76a]: routing procedure, flow control procedure, channel capacity assignment, priority queueing discipline, and topological configuration.

During implementation and system testing, analysis using simulated user input ensures that the network performs according to specifications, and isolates for reassessment any design flaws found.

Operational system studies are oriented toward system improvement by bottleneck detection, and capacity planning for expansion and upgrading of services. Measurements using real user traffic serve to validate and refine the system model developed during the network design stages.

The fundamental set of parameters for evaluation of resource sharing systems are: response time, throughput, resource capacity, and resource utilization.

Response time and throughput [KLEI76a] are closely related, but usually are applied to different kinds of traffic. Interactive traffic must be delivered quickly (fast response time). File transfers are concerned with how many bits per second can be pumped through the network (high throughput). A network may also be required to handle real time traffic which demands low delay and high throughput simultaneously.

Resource capacity is defined as the maximum rate at which a system can perform work. Resource utilization is the expected fraction of the system's capacity that is in use, reflecting the way in which system performance varies with the average system load.

In this thesis, computer network performance measurements and measurement techniques are reviewed. The thesis structure reflects the categorization and classification of the information contained in the referenced literature according to subject area. The sources of the data are referenced early in each section, and are also cross-referenced to other chapters and papers in the Bibliography. Tables are used extensively to summarize reported results, as available, and also to present the information in a standardized format for comparison purposes.

Chapter 2 describes a computerized database containing a comprehensive bibliography on computer network design, modelling, performance measurement and evaluation. From this database, the subset containing references to publications reporting network performance measurement data from real or simulated user input was selected. The subset, annotated and listed in the Bibliography, is the source material for the analysis presented in all of the following chapters.

Chapters 3 and 4 serve as an introduction to computer networks. Definitions, and a classification of the different types of networks are given. Activities in the area of international protocol standardization are discussed, and an overview is presented of the proposed model for process to process communication in distributed systems. The historical development and summary of the specific networks and their protocol implementations analyzed in later chapters is included. A list of other public network services available is also given.

Chapters 5 and 6 identify the goals and objectives of a network measurement activity, the requirements of a monitoring facility, and the hardware and/or software tools needed to collect the data. The types of statistics gathered, and the measurement objectives satisfied in different networks are compared.

A data flow model characterizing the user input process is presented in Chapter 7. Based on measurement experiments of real user traffic, observed values of the workload variables from different networks are noted.

Chapters 8, 9 and 10 are concerned with comparisons of network performance. The fundamental measures of throughput and delay in various networks are investigated, and the relative performance assessed. The efficiency of different protocols in managing the transfer of packets through a network is judged in terms of the amount of control overhead required. The reliability of networks in delivering messages accurately, and availability statistics on network connectivity are the final evaluation criteria.

Finally, the future of computer communication is predicted, and the role of network performance measurement placed in perspective.

2. Computer Network Design and Performance Database

A computerized database containing references to publications on computer network analysis and design, performance modelling, measurement and evaluation, compiled by the author, is described in this chapter.

The database, called NETWORKS, contains approximately 1000 entries. Each publication is fully catalogued, classified according to subject content, abstracted and keyworded.

The information is publicly accessible in printed form [TOWN80a], and as an online searchable database [TOWN80b]. A subset of the database, listed in the Annotated Bibliography, was selected for detailed analysis in the remainder of this thesis.

The database is an extension of a bibliography on computer system performance analysis [TOWN77] initiated by students of a graduate level course in Computing Science. The structure of the bibliography was enhanced when the NETWORKS database was implemented to provide more comprehensive retrieval of information.

In the sections below, the guidelines used for selecting documents for inclusion in the database are listed. A brief overview is given (Section 2.2) of the database management system used to implement the file, while further detail is given in Appendix A. Finally, the design features (record structure, searchable indexes, output formats) of the database are summarized (Section 2.3), with more detailed documentation given in Appendix B.

2.1 Guidelines for Document Selection

The following guidelines were established for determining whether or not a document is to be included in the database.

- Document Availability:

  The document is published and publicly available. The source of the papers is restricted to scientific and technical literature, for example, journals, conference proceedings, books, theses and reports. Private communications and internal reports are not included.

- Publication Language:

  The document is published in the English language. This includes translations, if available. Although not considered at this time, the database structure allows for entry of foreign language papers at a later date.

- Time Span of the Database:

  The time span covered is January, 1975, to November, 1979. Only key documents published prior to January, 1975, were added for the sake of completeness.

The database is comprehensive is subject coverage, but is not necessarily complete. Only those resources available from the University of Alberta were perused for selecting documents.

2.2 SPIRES Database Management System

The database management system (dbms) used for storage and retrieval of the NETWORKS bibliographic information is SPIRES (Stanford Public Information Retrieval System). SPIRES was designed and developed by Stanford University, and subsequently adapted for use at the University of Alberta Computing Center.

SPIRES was chosen for the NETWORKS database primarily because of its high level user interface and automated indexing feature. Optimized for online interactive retrieval, the powerful command language allows manipulation of search terms (index entries) using Boolean logic and set combination, truncated (stem) searching and range searching. Record items not indexed can be accessed using sequential search techniques.

The structure of a database is given to SPIRES in the form of a file definition. Included are rules for validating data on input and achieving data compression. Output rules translate codified data into more meaningful text strings. The structure of the indexes is declared, together with transformations necessary to create the indexes from the main record information. Access to the database by other computer accounts is also specified.

A report generating facility is achieved through formats. Formats specify, in a high level language, the mapping of record elements to a character matrix, according to a presentation layout. Different formats may be defined for online display and publication purposes.

Packaging of MTS (operating system), editor and SPIRES command sequences is possible via protocols. The protocol may prompt for user input: condition testing and branching capabilities are part of the protocol facility.

The SPIRES system is described in more detail with respect to physical file organization, privacy and security, database recovery, file management tools and documentation in Appendix A.

2.3 The NETWORKS Database

The NETWORKS database record items, searchable indexes, output formats, and protocols are now summarized. Additional details may be found in Appendix B; samples of records and output formats are contained in the User's Manual [TOWN80b].

The data items entered for each record in the database are:

- unique record number;
- primary author alphabetic and publication year identifier;
- date record added to database;
- date record was last updated;
- publication date;
- document type: one of book, report, journal, thesis or conference proceedings;
- title;
- authors;

- author affiliations;
- comments related to author, for example, research conducted while on leave of absence;
- journal name, volume and issue number for journal entries;
- publisher, location, edition number and Library of Congress number for book entries;
- conference name, theme, location, dates, proceedings editor(s), and sponsoring bodies for conference paper entries;
- publishing organization and report number for report entries;
- degree, faculty, university, and advisor for thesis entries;
- page numbers;
- level of detail in which the document was reviewed;
- broad classification (treatment) according to performance study area, type of paper, level of presentation, and type of data analyzed, from a vocabulary of 20 terms;
- keywords assigned to identify the subject content;
- abstract (author or annotated) and abstract number;
- section headings for journal or conference paper entries;
- table of contents for book, report or thesis entries;
- International Standard Serial Number;
- International Standard Book Number;

- document numbers of references cited;
- special notes for cross referencing purposes; and,
- language and publication country for foreign language entries.

Indexes, to facilitate online retrieval of references from the database, were constructed for:

- identifier;
- title;
- author;
- author affiliation;
- source;
- review depth, treatment, and document type;
- date;
- keywords;
- abstract; and,
- citations.

Three online display formats were written, each utilizing a 70 column output width suitable for video display terminals. The amount of record information printed ranges from the minimum citation, to the full record display.

The publishing formats produce double column output for the line/page printer. The listings are paginated, dated, and have a header at the top of each page to identify the source data.

The printed version of the NETWORKS database, called Computer Network Performance Bibliography [TOWN80a], contains the following sections:

- document citation, treatment words, subject keywords, and abstracts, printed in alphabetical order by identifier;
- alphabetical author list;
- author index;
- alphabetical author affiliation list;
- author affiliation index;
- alphabetical source list;
- source index;
- alphabetical keyword list;
- keyword index; and,
- alphabetical title list, with complete citation and document number.

The term lists give the number of publications attributed to the entry, and cross reference information to alternate index terms. The citations under each index entry are printed in reverse chronological order, i.e., most recently published documents first.

Two protocols were developed:

- to produce a printout of all the term lists; and,
- to produce the Computer Network Performance Bibliography.

2.4 Summary

The NETWORKS database is a useful tool for research in computer network design and performance analysis. Access to the literature is provided readily and easily.

3. Definitions and the Role of Standardization

In this chapter, the concepts central to packet transmission networks are defined. A brief description of the different types of networks follows, and finally, the role of international standardization is discussed.

3.1 Definitions

To establish the terminology framework for the remainder of the thesis, the following definitions [KLEI76a, MCQU77, KLEI78b] are offered. The definitions are sufficiently general as to be applicable to any network. The details with respect to specific network implementations are found in Chapter 4.

- Node:

  A node is a computer, with limited storage and processing resources, which directs the flow of packets within a network according to a specified protocol.
- Channel:

  A channel is a data transmission medium of relatively high bandwidth and low error rates. Examples are coaxial cables, radio channels, satellite channels and optic fibers.
- Host:

  A host is a computer, typically, but not necessarily, a time sharing system, connected to a node that is the source and/or destination of packets. A host may also perform the functions of a node.

- Topology:

    The network topology describes the locations of the nodes, and their channel connectivity, i.e., the physical configuration of the network.

- Message:

    A message is a logical unit of data, for example, a line of text terminated by a carriage return, exchanged between a source and destination host.

- Packet:

    A packet is a unit of data (message) exchanged between nodes. Messages are enclosed by headers (bits) containing addressing and error control information. One packet may contain many short messages, or long packets may be split into several packets.

- Acknowledgment:

    An acknowledgment is a control packet returned to the source from the destination to indicate successful receipt of the packet. Retransmission procedures are invoked if packet errors are detected, or the packet fails to arrive at its destination.

- Routing Algorithm:

    The routing algorithm is a procedure to determine which of several paths through a network will be taken by a packet enroute to its destination.

- Flow Control:

    Flow control is an algorithm to regulate the rate at which data enters and leaves the network. It is a mechanism to match the source input rate to the destination acceptance rate by limiting the number of unacknowledged packets in transit.

- Sequencing and Numbering:

  Packet headers may contain control information to ensure delivery of packets in the same order as received. Duplicate and lost packets can thus be detected. Not all networks provide this feature.

- Error Control:

  Control information contained in packet headers is used to detect bit errors resulting from channel noise.

- Protocol:

  A protocol is a collection of acknowledgment, routing, flow control, sequencing and numbering, and error control rules to govern packet transfer in a network. The protocol is executed by the nodes.

There are many variations in the realization of these concepts: see Chapter 4 for some specific examples of network implementations.

3.2 Types of Networks

There are two fundamental approaches to the arrangement and specification of the basic units of a packet transmission network. They differ in terms of how the channel is used, and where in the system queueing of packets occurs.

The two classes of packet transmission networks are:

- packet switched; and,
- multiaccess.

Definitions of each are given below, together with the salient features of each category.

- Packet Switched Networks:

  A packet switched network generally consists of a collection of nodes connected together in point to point fashion by high speed transmission channels (typically digital trunk lines). The nodes are often located more than one kilometer apart: the network topology usually does not form a fully connected graph, but more than one path almost always exists between any two nodes on the network.

  Packets queue at the entry points to the network. Once in the network, the packets are transferred from node to node until they arrive at their destination site. A packet is retained at a node (nodal queueing delay) until the outgoing transmission line ceases to be busy, and buffers are available at the next node in the path. Thus, optimization of both switch and channel capacity, together with network topology, are the goals in network design to achieve high utilization without excessive delay.

- Multiaccess Networks:

  Multiaccess networks typically utilize only a single channel, and differ from packet switched networks in the way the channel is used. All nodes transmit into the channel (multiaccess), and receive all messages sent (broadcast) on the channel. The node contains sufficient intelligence to accept only those packets addressed to it.

Thus, queueing occurs at the entrance to the network, but not within the system itself. Network design concentrates on optimizing the channel access scheme by minimizing the number of collisions (overlaps in packet transmissions, causing both packets to be rejected).

Three different types of communication channels may be used, depending on the intended function of the network. Satellite channels are used for long distance transmissions, such as broadcasting overseas or to remote locations as, for example, the far North or developing countries. Radio channels are used over medium range distances, predominantly for local distribution of messages. Coaxial cables, less than one kilometer in length, are typically used for local area network applications.

3.3 Open System Interconnection

In this section, the specification for an international protocol standard based on packet switched services is outlined.[1] The purpose is to establish a framework for the next chapter, in which some specific packet switched networks are described.

As yet, there are no proposals for a standard multiaccess protocol.

In the fall of 1976, the International Telephone and Telegraph Consultative Committee (CCITT) approved the X.25 Recommendation for interfacing hosts and terminals to packet switched networks.

A subcommittee of the International Standards Organization (ISO) was formed in the spring of 1977 to extend the X.25 Recommendation [BACH78a, BACH78b, BLAC78, DESJ78, EMMO78, MACD78] to include higher level protocols for process to process communication in distributed systems. A seven layer architecture, known as the Open System Interconnection (OSI), is the proposed structure: the X.25 Recommendation encompasses the first three (lowest) layers.

Vendor supplied distributed system products allow interconnection of that manufacturer's terminals and computers, but do not interface easily with other vendor product lines. The OSI is a formalization of functions, rather than a specification for a protocol implementation procedure.

Communication in OSI is factored into simple tasks, which are organized into independent layers of functionality. Each layer implements a set of services accessible only by the layer immediately above.

The first four layers are named the "transport service" and provide the mechanism to transport messages from one workstation (computer system, intelligent terminal, etc.) to another via a network (packet switched or multiaccess), or several interconnected networks.

The upper three layers are called "users of transport service" and control the interaction between processes.

Figure 1 illustrates two processes communicating within a workstation, and two processes communicating between workstations. The flow of intranetwork messages through switching nodes, and internetwork messages through network gateways, in the transport service is shown.

The function of each of the seven layers is now described, in order of the lowest hierarchical level to the highest.

- Physical Control Level:

  The function of the physical control level is to provide a means of bit transmission across a physical link (channel), independent of medium.

- Link Control Level:

  The link control level enables messages (series of bits) to be exchanged reliably across a single physical link (between nodes). Error detection and correction, link sequencing and flow control, and retransmission procedures are all functions of the link control level.

- Network Control Level:

  Management of intranetwork operations is the function of the network control level. Addressing, routing, message priority control, flow control,

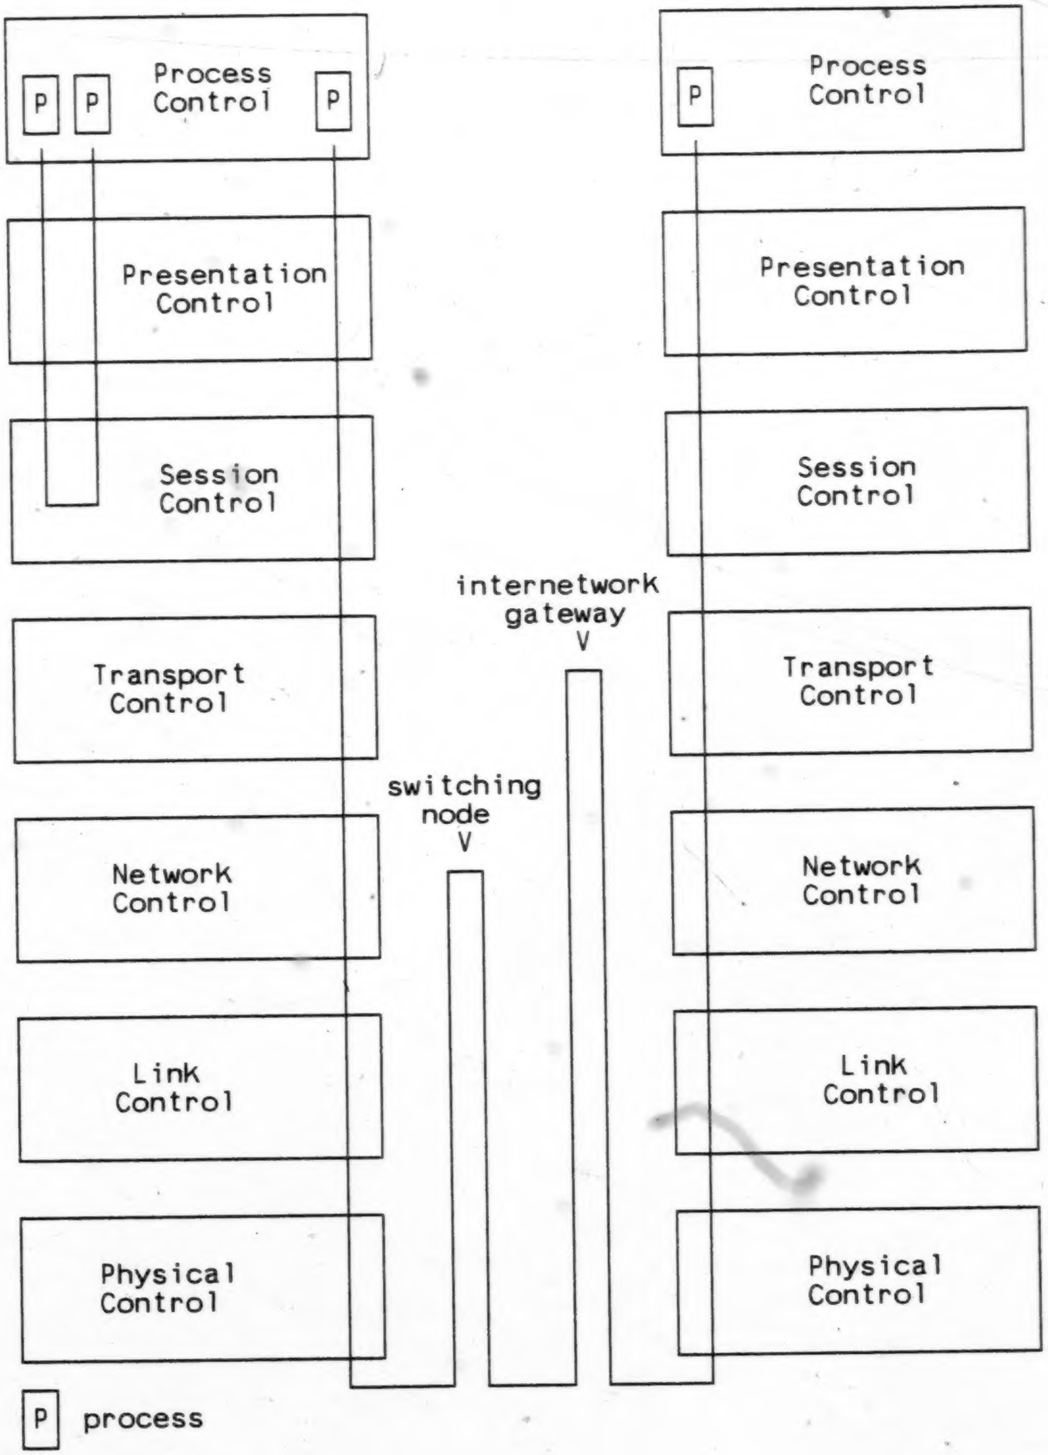
Process
Control
P
P
P
Presentation
Control
Session
Control
Transport
Control
Network
Control
Link
Control
Physical
Control
internetwork
gateway
V
switching
node
V
P
process

Figure 1 - Process to Process Communication Within and Between Workstations in Open System Interconnection

security and accounting are duties of this level.

- Transport Control Level:

  Concerned with end to end control of message transfers, the transport control level directs the flow of packets across an arbitrary network topology, including several interconnected networks.

- Session Control Level:

  The role of the session control level is to establish, maintain and terminate logical connections between processes. Synchronization of message transfers, security, and journalization of information exchanged to support integrity requirements are all provided by the session control level.

- Presentation Control Level:

  Device handling, data code transformations and encryption are all functions handled by the presentation control level.

- Application (or Process) Control Level:

  The user applications, included for the sake of completeness, make up the process control level.

The OSI is not a completely new concept, rather it is based on currently accepted international standards and procedures. In Figure 2, the relationship and placement of OSI functional levels to existing standards is shown. The expansion of the acronyms and abbreviations used are given in Table 1, as well as other protocol proposals not yet formally accepted.

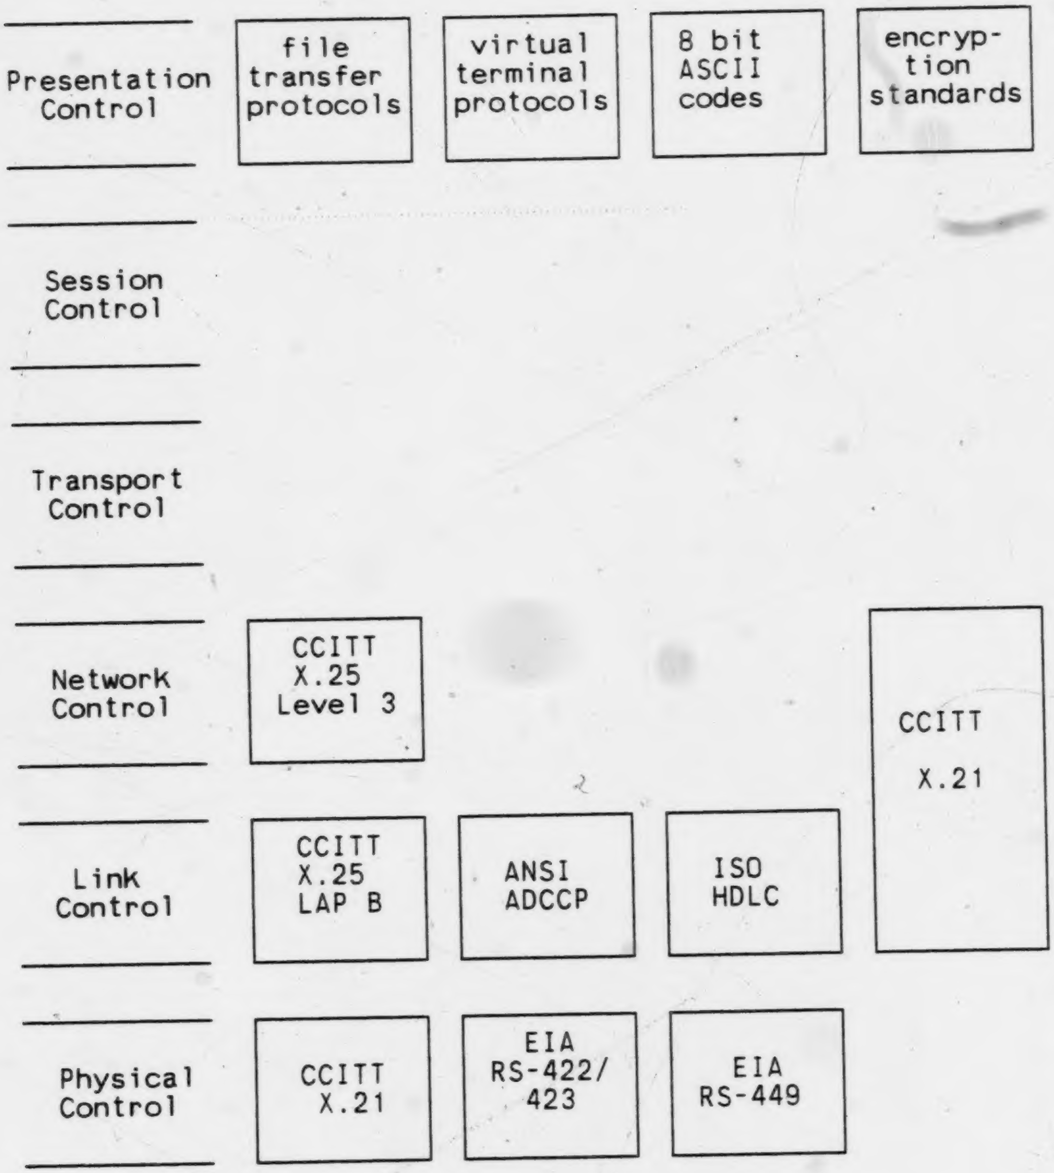
Presentation Control
file transfer protocols
virtual terminal protocols
8 bit ASCII codes
encryption standards
Session Control
Transport Control
Network Control
CCITT X.25 Level 3
CCITT X.21
Link Control
CCITT X.25 LAP B
ANSI ADCCP
ISO HDLC
Physical Control
CCITT X.21
EIA RS-422/423
EIA RS-449

Figure 2 - Accepted International Standards in Relation to Open System Interconnection

Level 1	
CCITT X.21	General Purpose Interface Between Data Terminal Equipment (DTE) and Data Circuit Terminating Equipment (DCE) for Synchronous Operation on Public Data Networks
EIA RS-422/423	Electrical Characteristics of Balanced/Unbalanced Voltage Digital Interface Circuits
EIA RS-449	General Purpose 37-position and 9-position Interface for Data Terminal Equipment (DTE) and Data Circuit Terminating Equipment (DCE) Employing Serial Binary Data Interchange
Level 2	
CCITT X.25 LAP B	X.25 Link Access Protocol B
ANSI ADCCP	Advanced Data Communication Control Procedure (BSR X3.66)
ISO HDLC	High Level Data Link Control
FTS 1003	United States Government proposed Federal Telecommunication Standard 1003
IBM SDLC	Synchronous Data Link Control

Table 1 - International Communication Standards and Activities

Level 3

CCITT X.25	Interface standard between a network and a single endpoint
CCITT X.21	Standard for circuit switched networks defines functions for both level 2 and 3
ISO/TC97/SC6 Project 24	Subcommittee working on network control protocols between peer endpoints
ANSI/X3 Project 281	Same as above

Level 4 and above

ANSI/TG X3S33	May combine the functions of Levels 3 and 4
CCITT SGVII	Studying standards for transport end to end control
ISO/TC97/SC16 and SC6	Studying standards for Open System Interconnection

Table 1 (cont'd) - International Communication Standards and Activities

The OSI specifications are expected to be completed by December, 1980 [BACH78a]; international review and acceptance will most likely follow.

3.4 Summary

Identified and defined were the basic concepts of packet transmission networks, including the characterization of the two fundamental types of networks: packet switched and multiaccess. The proposed standard for process to process communication in a packet switched network environment was presented.

4. Network Protocols and Services

Having established a reference point for packet transmission networks in the previous chapter, an overview is now presented of specific network transport service implementations, and the protocols that they employ.

The descriptions are presented at a fairly high level: a formal specification for analytic comparison was not intended. The objective is to provide a basis for understanding the performance measurement experiments discussed in later chapters.

4.1 Packet Switched Networks

Packet switched networks were first proposed in the late 1960's [ROBE78]. The earliest networks (ARPANET and NPL, see Table 2) demonstrated the feasibility and viability of the technology.

As more countries and organizations became interested in providing packet switched network services, the set of standards (CCITT X.25) for interfacing computers and terminals to a network was developed, and endorsed, in the fall of 1976. Computer networks utilizing the X.25 Recommendation have the added advantage of compatibility for global interconnection of services.

Table 2 summarizes the non-X.25 packet switched networks; Table 3 lists the X.25 compatible services. The corporate sponsors of each project are given, together with

	Sponsor(s)	Date Operational	Configuration
ARPANET (Advanced Research Projects Agency)	U.S. Dept. of Defence; Bolt Beranek and Newman, Inc.; UCLA	1969 [ROBE78, KLEI76a]	200+ nodes
CTNE Network (Compania Telefonica Nacional de Espana)	Spanish PTT and group of users	1973 [POUZ79]	2 nodes plus a number of concentrators
CYCLADES	French government and universities	1974 [ROBE78, POUZ79]	17 nodes, 9 concentrators, 18 hosts
EIN Network (European Informatics Network)	multination European research network	1976 [ROBE78]	unknown
EPSS Network (Experimental Packet Switched Service)	British Post Office, X.25 service expected in 1980 or 1981	1977 [POUZ79, ROBE78]	3 nodes
General Electric International Network	General Electric (probably Honeywell now)	[WEDB74]	330+ entry points
INFOSWITCH Network	CNCP Telecommunications, Canada	[POUZ79]	unknown

Table 2 - Non X.25 Public Packet Switched Networks

	Sponsor(s)	Date Operational	Configuration
KUIPNET	Kyoto University, Japan	1977 [KITA78]	1 node, 6 hosts
NPL Network (National Physical Laboratory)	National Physical Laboratory, England	1973 [ROBE78]	1 node
TYMNET Network	Tymshare, Inc., United States	1971 [ROBE78, POUZ79]	300+ nodes

Table 2 (cont'd) - Non X.25 Public Packet Switched Networks the commercial service date and network configuration, if known.

More specific details of networks whose performance is analyzed in later chapters are now given.

4.1.1 ARPANET Network

ARPANET, under the direction of the Advanced Research Projects Agency (ARPA) of the United States, was the first packet switched network. Planning [FRAN72, MCKE72, KLEI74, KLEI76a, MCQU77, ROBE78] began in 1967, with the first four nodes operational in December, 1969. The network expanded to 23 hosts by April, 1971, 62 hosts by June, 1974, and 111 hosts by March, 1977.

Bolt Beranek and Newman, Inc., was responsible for the design of the packet switching equipment and development of the network software and control functions. Theoretical analysis and performance measurement work was carried out by the University of California at Los Angeles (UCLA).

Minicomputers, called Interface Message Processors (IMPs), are used as node switches. The IMPs are connected together by 50 Kbps (kilobits per second) digital trunk lines.

Messages [MCKE72, KLEI74, KLEI76a] are subdivided into small fixed length units called packets: up to 8 packets can be created for each message. The packets are transferred individually from IMP to IMP using an adaptive routing scheme. Packets are acknowledged between IMP pairs, and

	Sponsor(s)	Date Operational	Configuration
DATAPAC	Trans Canada Telephone System; Bell-Northern Research	1977 [MCGI78]	55 nodes
Dutch PTT Network	Dutch PTT	1980 [POUZ79]	unknown
DX-2 Network	Nippon Telephone and Telegraph Co. Ltd., Japan	1979 [ROBE78]	unknown
EURONET Network	9 European Common Market countries	1979 [ROBE78, POUZ79]	4 nodes and 3 concentrators: is a copy of TRANSPAC
IBM Experimental Satellite Link Network	IBM Corporation	1977 [ULME78]	2 nodes, and satellite channel
RCP Network (Reseau a Commutation par Pacquets)	French PTT	1974 [ROBE78]	testbed for TRANSPAC
TRANSPAC Network	French PTT	1979 [ROBE78, POUZ79]	12 nodes and 13 concentrators

Table 3 - X.25 Compatible Public Packet Switched Networks

	Sponsor(s)	Date Operational	Configuration
TELENET Network	Telenet Communications Corporation, United States	1975 [ROBE78, POUZ79]	200+ nodes

Table 3 (cont'd) - X.25 Compatible Public Packet Switched Networks reassembled at the destination IMP prior to delivery to the host. The host then returns a "Request for Next Message" (RFNM) control packet which serves as both flow control and an end to end acknowledgment.

ARPANET was the first network to utilize a completely distributed, dynamic routing algorithm.

4.1.2 CYCLADES Network

CYCLADES [POUZ74, LELA76, EYRI77, GELE77, GRAN77, GIEN78, ROBE78, POUZ79] is a computer network linking research centers and universities in France. Funded by government sponsorship, CYCLADES has been operational since 1974. At the end of 1978, the network configuration consisted of 7 nodes, 9 concentrators and 18 hosts.

The packet switched communication network within CYCLADES is called CIGALE. CIGALE is only a datagram service, simply sending blocks of data across the network without guaranteeing delivery or orderly flow of packets. Packet sequencing and duplicate detection are the responsibility of the host. There is no flow control mechanism, but alternate path routing has been implemented.

CIGALE uses one of two protocols between nodes.

- TMM Protocol:

  TMM uses "stop and wait" acknowledgment control. An acknowledgment must be received prior to sending the next packet.

- MV8 Protocol:

  The transmission line is considered to have 8 virtual channels. Packets are sent on free channels, thus up to 8 packets may be in transit between nodes. Three variations of the channel allocation scheme are available: fixed and dynamic allocation based on packet length, and FIFO.

4.1.3 DATAPAC Network

The DATAPAC network [CLIP76, COHE78, MCGI78, ROBE78, POUZ79] is a Canadian, public, packet switched data communication network operated by the Trans Canada Telephone System (TCTS).

DATAPAC, developed by Bell-Northern Research for TCTS, was an X.25 based network from the start. DATAPAC began commercial service on July 15, 1977. By August, 1979, DATAPAC nodes were located in 55 centers across Canada. The nodes are linked by high speed digital trunk lines.

Hosts and intelligent terminals interface to DATAPAC using the Standard Network Access Protocol (SNAP). Asynchronous, unbuffered terminals may access a node via a Network Interface Machine (NIM) providing Interactive Terminal Interface (ITI) services. Interface support is also provided for controllers handling multidrop terminals. Unique features of DATAPAC service include two levels of message priority, collect call schemes, and closed user groups.

DATAPAC is interconnected with the TELENET and TYMNET networks of the United States. An experimental link to the French RCP network has been demonstrated.

4.1.4 IBM Experimental Satellite Link Network

In 1977, an experimental network [ULME78] was established to test the performance characteristics of the point to point High Level Data Link Control (HDLC) protocol. The network configuration consisted of two IBM 370/158 computers located in Gaithersburg, Maryland, and La Gaude, France, linked by a Symphonie 1.544 Mbps satellite channel. The Symphonie satellite was at that time owned by a French-German consortium.

4.1.5 KUIPNET Network

The Kyoto University Information Processing Network (KUIPNET) is a local area, star network configuration [KITA78] connecting six heterogeneous computers to a single switching computer. Speech data, digital images, file and interactive traffic types are all handled by the network.

KUIPNET has a gateway to the Japanese Experimental Network for Inter-University Data Processing Centers.

4.2 Multiaccess Networks

Implementations of multiaccess networks are presented in Table 4. Listed are the corporate sponsors, service dates and network configuration, if available.

Three types of communication channels are being tested for multiaccess networks.

- Radio Channels:

  The study of radio communication networks is an investigation into the replacement of point to point (host to terminal) wire connections by radio channels. Additionally, radio based networks are particularly well suited for mobile terminal to computer communication. The two networks utilizing radio channels are ALOHANET and PRNET.

- Satellite Channels:

  Satellite channels are characterized by high bandwidth with a fixed propagation delay (approximately .25 seconds) which is large relative to the packet transmission time. Because of the wide area of coverage offered by satellite broadcasts, access to computer facilities by terminals in remote locations is possible. The multiaccess network utilizing a satellite channel is SATNET.

- Cable Channels:

  Local area computer network applications utilize multiaccess schemes with a wire based communication line. Distances are limited to less than one kilometer

	Sponsor(s)	Date Operational	Configuration
ALOHANET	University of Hawaii	1971 [KAHN78]	two radio channel, star configuration
ETHERNET	Xerox Palo Alto Research Center	1976 [METC76 SHOC79b]	coaxial cable channel, 120 hosts
FIBERNET	Xerox Palo Alto Research Center	1978 [RAWS78]	optic fiber channel
PRNET	Advanced Res. Projects Agency; Bolt Beranek and Newman, Inc; SRI International; Rockwell International; UCLA; Network Analysis Corp.	1978 [KAHN78, KUNZ78, TOBA78]	radio channels, 2 stations, 4 repeaters, 4-5 mobile terminals
SATNET	Advanced Res. Projects Agency; Defense Communications Agency; Norwegian Telecommunications Administration; British Post Office	1975 [BRES79]	satellite channel, 4 earth stations

Table 4 - Multiaccess Networks

of cable. Local area networks have been described [METC76, CLAR78, RAWS78, SHOC79a] as the midpoint between the extremes of remote computer networking, as exemplified by ARPANET, and multiprocessor architectures employing parallel processing. The ETHERNET network is an example of a local area network utilizing multiaccess channel techniques.

Specific details of multiaccess networks whose performance is analyzed in later chapters are now given.

4.2.1 ALOHANET Network

The ALOHANET packet broadcasting radio network [ABRN73, BIND75, KAHN78] has been in operation at the University of Hawaii since 1971. ALOHANET was the first system to utilize the packet broadcasting concept for online access of a central computer via radio.

The central communications processor is called the MENEHUNE (Hawaiian for IMP) and forms the center of a two channel, star network configuration.

Messages, up to 86 bytes in length, are sent to the MENEHUNE via the remote access radio channel using the pure ALOHA protocol. Computer to user messages and packet acknowledgment traffic is sent on a separate broadcast channel. Repeaters, which act as simple store and forward devices, extend the range of radio communication.

The pure ALOHA protocol [ABRN73, KLEI73, BIND75, LAM77, JACO78, KAHN78] is a random access scheme in which packets are sent to the node in a completely unsynchronized manner. If two packets are broadcast at the same time, or their transmissions overlap, then a collision occurs with both packets being rejected.

If an acknowledgement is not received at the terminal within a specified time period, retransmission of the packet occurs after a randomized delay.

The broadcast channel uses priority queueing to ensure that acknowledgment traffic is sent before user data packets. Only one packet per user is allowed in the buffer pool, and service ordering is based on FIFO.

ALOHANET is connected to ARPANET in California by an INTELSAT IV 50 Kbps satellite channel. Extensions were made to MENEHUNE to accommodate the ARPANET flow control, multiline packets and full duplex mechanisms.

4.2.2 ETHERNET Network

The ETHERNET [METC76, SHOC79b] is a local area network designed and implemented by the Xerox Palo Alto Research Center. The ETHERNET measurement experiments [SHOC79b] were carried out on a configuration of over 120 machines on an 1800 foot 2.94 Mbps coaxial cable.

ETHERNET utilizes the Carrier Sense Multiple Access (CSMA) with collision detection protocol [METC76, SHOC79b], which is based on the pure ALOHA protocol (Section 4.2.1). Prior to sending a packet onto the channel, a carrier sense (hardware) mechanism is used to determine if a transmission is already in progress. If the channel is free, the packet is sent immediately; otherwise the station waits until the transmission is complete.

It is still possible to have collisions if two or more stations sense a free channel and transmit simultaneously. Each station monitors the channel, and detects a collision when the cable signal does not match its own transmission. Retransmission occurs after a wait time of random length: the mean retransmission interval is adjusted by the station in proportion to the frequency of collisions.

4.2.3 Packet Radio Experimental (PRNET) Network

Commissioned by the Advanced Research Projects Agency (ARPA), an experimental network using a radio channel [TOBA76, KAHN78, KUNZ78, TOBA78] is being assembled and tested by the SRI Telecommunications Sciences Center. The hardware was designed by the Collins Radio Group of Rockwell International; software development was carried out by Bolt Beranek and Newman, Inc., and SRI International. Other contractors in the project are: Network Analysis Corporation, and University of California at Los Angeles.

PRNET is the first network to combine multihop routing capabilities with multiaccess shared radio channel technology. There are three types of components:

- packet radio terminals, which are the sources and destinations of traffic;
- packet radio repeaters, which extend the broadcast

range by acting as store and forward relays; and,

- packet radio stations, which provide global network control.

The channel access protocol is either pure ALOHA, or CSMA (Sections 4.2.1 and 4.2.2). Routing in the network is based on a tree configuration determined by the radio distance (number of hops) from the station.

Two experimental implementations of PRNET are operational:

- San Francisco Bay Area Packet Radio Network, consisting of 4-6 mobile terminals, 4 repeaters, and 2 stations; and,
- Packet Radio Laboratory Testbed, consisting of a group of packet radios that can simulate various connectivity arrangements, and a minicomputer based instrumentation system for testing radio channels.

PRNET has an internetwork gateway to ARPANET implemented in the packet radio station.

4.2.4 SATNET Network

The SATNET (Atlantic Broadcast Satellite Network) is an experimental packet broadcast network sponsored by the Advanced Research Projects Agency (ARPA) and Defense Communications Agency (DCA) of the United States, the British Post Office, and the Norwegian Telecommunications Administration [GERL77b, JACO77, CHU78, JACO78, KLEI78a, TREA78, WEIS78, BRES79].

The experiment was initiated in mid-1975 to investigate the feasibility and efficiency of different channel access schemes, the integration of speech and data traffic, voice conferencing, and host SATNET protocols. Participants in the project are: LINKABIT Corporation, Bolt Beranek and Newman, Inc., University of California at Los Angeles, Massachusetts Institute of Technology Lincoln Laboratory, and Communications Satellite Corporation (COMSAT), all of the United States; University College London, England, and the Norwegian Defense Research Establishment, Norway.

The configuration of the network consists of a 64 Kbps INTELSAT IV-A SPADE satellite channel, and four earth stations located in Etam, West Virginia, Clarksburg, Maryland, Goonhilly Downs, England, and Tanum, Sweden. The ground stations are called SIMPs (Satellite IMPs) and are modelled after the IMPs of the ARPANET.

Five different protocols were tested in the SATNET experiments:

- Fixed-TDMA (F-TDMA) protocol;
- Reservation-TDMA (R-TDMA) protocol;
- Slotted-ALOHA (S-ALOHA) protocol;
- Controlled S-ALOHA protocol; and,
- Contention Based Priority Oriented Demand Assignment (CPODA) protocol.

For the first four protocols, the channel is logically subdivided into time slots, sufficiently long to transmit one maximum size packet. Slots are then grouped into frames

The fifth protocol uses a fixed frame size, but varies the slot length according to packet size. A frame size of 32 slots was used for all SATNET experiments.

More specific details are now given for the protocols.

- Fixed-TDMA (F-TDMA) Protocol:

  F-TDMA is a static reservation system [ROBE73, GERL77b, KLEI78a, WEIS78] wherein the channel slots are preallocated and equally subdivided among all stations. A station transmits only during its allocation, and cannot dynamically request more channel capacity. Unused slots are not reassigned to other users.

- Reservation-TDMA (R-TDMA) Protocol:

  The R-TDMA protocol [GERL77b, LAM77, JACO78, KLEI78a, WEIS78] is a demand access scheme. Like F-TDMA, channel slots are assigned to stations. However, unused slots may be allocated to stations with packets to send.

  The first slots in a frame are called reservation subframes: each station uses their subframe to indicate the number of packets they have awaiting transmission. The channel scheduler, a distributed algorithm running synchronously in all stations, then assigns the unused slots on a round robin basis. If noise prevents the station from receiving all reservations correctly, the scheduler reverts to the F-TDMA protocol for the remainder of the frame.

- Slotted-ALOHA (S-ALOHA) Protocol:

In slotted-ALOHA [ABRN73, KLEI73, ROBE73, GERL77b, KLEI77, LAM77, JACO78, KAHN78, KLEI78a, TOBA78, WEIS78], a random access scheme, stations transmit packets (beginning on slot boundaries) whenever they are ready. If two or more stations transmit in the same slot, the packets collide, and require retransmission.

Each station maintains two output queues: a new queue for new packets, and a retransmit queue for packets requiring retransmission. The scheduling algorithm first transmits a packet from the retransmit queue with probability PR. If the retransmit queue is empty, then the station transmits from the new queue with probability PN. Packets arriving at an empty station are transmitted in the next possible slot (probability P = 1). The service ordering in each queue is FIFO.

- Controlled Slotted-ALOHA Protocol:

The controlled S-ALOHA protocol [KLEI77, JACO78, KLEI78a, TOBA78, WEIS78] is based on S-ALOHA above, but adjusts the PR and PN gate values dynamically according to changing traffic patterns. The proposal is based on the observation [ABRN73] that for arbitrary numbers of stations and traffic patterns, maximum throughput is obtained when the average number of packets transmitted per slot is equal to 1.

The PN and PR values for each station are dynamically adjusted based on the collision rate and channel load, estimated over a window (number) of slots. In heavy traffic, the controls maintain equilibrium, i.e., PN(i) = PN(j), for all stations i, j, and the summation of all PN(i) = 1, thus achieving optimal channel utilization.

- Contention Based Priority Oriented Demand Assignment (CPODA) Protocol:

The CPODA protocol [JACO77, CHU78, JACO78] is designed to satisfy more sophisticated traffic requirements. The protocol handles:

- an arbitrary mix of bursty and nonbursty (data and voice) traffic types;
- different delay classes and message priorities;
- a widely varying traffic load (input rates);
- messages that are sent to single or multiple destinations; and,
- an arbitrary mix of channel receiving rates for earth stations.

The channel frame is subdivided into two parts: a reservation subframe and an information (messages) subframe.

Reservation packets are submitted on a contention basis (using S-ALOHA) during the reservation subframe. The size of the reservation subframe grows and shrinks dynamically according to demand within the constraints of minimum and maximum size. Reservations may also be piggybacked in the headers of messages.

Service ordering of message transmissions is done according to a weighted function of priority and delay class. A distributed algorithm running in each earth station determines the channel scheduling. The messages are scheduled for consecutive transmission to minimize unused channel bandwidth, in contrast to protocols that wait for slot boundaries to begin transmission.

Packet acknowledgments are indicated either by piggybacking (up to 5 acknowledgments) in a message header or reservation request packet, or by sending an explicit acknowledgment message packet.

SATNET and ARPANET are connected together by gateway minicomputers. The gateways allow SATNET to operate independently, and also provide a convenient testbed for internetwork experiments.

4.3 Summary

This chapter has served to place in perspective the relative protocol complexities, service functions, distance ranges and stages of development of various network implementations. The performance of the different types of networks can be expected to vary accordingly, and this aspect will be discussed in later chapters.

The measurement experiments discussed in later chapters assume knowledge of this chapter.

5. Network Measurement Planning

In this chapter, the motivation for network measurement is developed by defining the goals and objectives of measurement experiments. The requirements of monitoring facilities are listed, and finally, the types of data collected are broadly categorized.

5.1 Network Measurement Goals and Objectives

The function of a network is to transfer messages from one location to another. Intuitively, network users expect the transfer to occur rapidly and accurately, at minimum cost and maximum convenience.

The process of analyzing how well a network executes its intended functions is called performance evaluation. The main criteria [KLEI76a] used to evaluate network performance are:

- throughput;
- delay;
- cost; and,
- reliability.

Throughput is defined as the rate at which data can be transferred through the network (Chapter 8, Sections 9.1 and 9.2). Delay is the time interval from the start of message transmission until it successfully arrives at its destination (Chapter 8). Cost is measured as the amount of network control overhead required to manage the packet transfer (Section 9.3). Reliability assesses the packet error detection and correction capabilities, as well as the rate of node and channel hardware failures (Chapter 10).

All of these performance criteria can be measured for a given user input load and network implementation. This is the type of performance data analyzed in later chapters.

For network designers, the behavior of a system under variations in load and implementation, is of greater interest. For example, the following questions might be asked:

- what happens when the user input characteristics change?
- what happens when protocol parameters are varied?
- what happens when the protocol algorithms are changed? or,
- what happens when the node and/or channel hardware is upgraded?

Investigation of solutions to these problems fall into the areas of performance projection and capacity planning. Performance projection predicts the effects of increased load, optimized algorithms and/or hardware reconfiguration on system behavior. Capacity planning evaluates the effects of hardware upgrades and/or replacements.

Although performance projection and capacity planning are not considered in further detail, network measurement plays an important role in these two activities.

More specifically, the goals of measurement for network performance evaluation, projection and capacity planning [COLE72, MORG74, WEDB74, TOBA76, GERL77b, TOBA78, GREE79] can be given as:

- software certification;
- model validation and refinement;
- user behavior characterization;
- system behavior analysis;
- system parameter optimization;
- efficient resource allocation;
- availability and reliability maximization; and,
- provision of a charge mechanism.

The objectives are discussed in more detail below, with references to later chapters or cited literature.

- Software Certification:

  Measurements are used to certify that system software is implemented correctly and performing according to specifications based on analytic model predictions. Testing and debugging of software usually requires simulated user input, and may involve the use of specialized tools to assist in detection of implementation errors. Examples of such specialized tools are the protocol testers developed for DATAPAC and TRANSPAC [BLEV78, WEIR78] to validate a host to network interface.

- Model Validation and Refinement:

  In modelling and simulation studies, simplifying assumptions are often made for model tractibility and/or ease of computation. Measurement experiments allow network designers to observe the behavior of the system under real traffic and a range of simulated traffic conditions. Experimental results are then used as a feedback mechanism for model validation and refinement.

For example, assumptions are made with respect to the user input process, e.g., average message and packet size, when selecting the optimal buffer size for the nodes. Real user traffic measurements (Section 7.2) revealed message sizes much lower than the estimated sizes: poor buffer utilization resulted.

Delay models are usually given with delay measurement data (Chapter 8), and are generally found to underestimate the actual delay due to factors not included in the model.

- User Behavior Characterization:

The user input process is independent of the network; the network behavior, however, depends on the nature of the input process.

Assumptions about the user workload, e.g., packet and message size distribution, rate of arrival of packets and messages, and traffic distribution, are central to network design. For example, traffic is assumed uniformly distributed throughout a network. In fact, preferred hosts and routes prevail (Sections 7.2.3 and 7.2.4).

- System Behavior Analysis:

  The behavior of a system is measured to derive observed values of performance evaluation parameters.

  Implementation inefficiencies and/or design flaws are suggested by unusual results or unanticipated network behavior. For example, improper synchronization in the updating of the parameters of the ARPANET distributed routing algorithm [KLEI76a, KLEI77, KLEI78b] led to "loop traps". A packet was repeatedly shuttled between two nodes, until the resulting congestion caused the packet to take an alternate path.

- System Parameter Optimization:

  System parameters are adjusted (fine tuned) to optimize network efficiency subject to a given user input load. Parameter estimates may be derived analytically: measurement results provide the "best" value to use.

  For the CPODA protocol (Section 4.2.4), measurement studies were conducted [CHU78] to select the optimal frame size, and minimum reservation subframe size. Similarly, the effects of varying the window size and probability increment size in the controlled S-ALOHA protocol (Section 4.2.4) have been studied [KLEI78a].

- Efficient Resource Allocation:

  The balance between network capacity and user demand is monitored to assess system efficiency. The objective is to achieve high throughput and low delay without wasting excess capacity.

The detection of bottlenecks can lead to changes in hardware and/or software configuration. As an example, an internetworking project, where PRNET is used as a linkage network between two ETHERNETS, showed that throughput was limited (Section 8.3) by the processing speed of the packet radio units.

- Availability and Reliability Maximization:

The detection of failsoft errors allows implementation of corrective procedures (Section 10.1). For example, packet bit errors and lost packets cause retransmission of the packet by the source node. Duplicate packets are discarded at the destination.

Hardware failure of a network component eliminates the node or channel connection to the network (Section 10.2). Routing tables must be updated to reflect the new network configuration.

- Provision of a Charge Mechanism:

Statistics collected on network usage are processed and charged to the user to recover costs of providing network services. However, charging algorithms and rate schedules are beyond the scope of this thesis.

5.2 Requirements of a Measurement Facility

The general requirements [MORG74, GREE79] of a network measurement facility are:

- to observe and record parameters of system behavior;
- to offer a choice in the resolution in the unit of measure;
- to be flexible in the selection of measures collected;
- to be system independent;
- to be easy to use;
- to operate with a minimum amount of human intervention;
- to minimize the amount of network interference, i.e., to be non-biasing;
- to have no effect on the security and integrity of the measured system;
- to provide simulated user input for controlled monitoring of system behavior;
- to provide meaningful statistical summaries of collected data; and,
- to be low cost, and efficient in operation.

Design of monitoring systems satisfying the above criteria is nontrivial: examples of implementations are presented in the next chapter.

5.3 Measurement Data Types

The types of measurement data [MORG74, GREE79] collected and their relevance to performance goals and objectives, can be classified into 5 broad categories:

- time measurements, for response time and delay analysis;
- space and time measurements, for resource utilization analysis;
- event count measurements, for input load analysis;
- length measurements, for resource capacity analysis; and,
- error measurements, for reliability analysis.

The measurements are used to gain insight into network behavior. In some cases, the probability distribution function of the parameters are analyzed, in others, a sample of data values is sufficient. For example, the average message and packet size are insufficient to characterize all of the user data traffic types (Section 6.1). The distribution function of message length for ETHERNET (Section 7.2.1) reveals a bimodal curve, corresponding to the short interactive and long file transfer data types.

The distribution function may be examined over the network as a whole, or for an individual network component, depending on the performance measure(s) of interest. For some measures, both may be analyzed. In multiaccess network studies, not only is the throughput of the entire network assessed (Sections 8.2, 9.1 and 9.2), but also that of each individual station in order to evaluate the fairness of the channel allocation scheme, i.e., according to each station's requirements.

5.4 Summary

The role of measurement in network performance evaluation, projection and capacity planning is evaluated. Specific examples of how measurements met performance goals and objectives are given. The desirable features of a measurement facility are listed, and finally the types of data collected are classified.

6. Network Measurement Tools

In this chapter, some representative implementations of network measurement facilities are evaluated. From the analysis, it is possible to classify the different types of monitors, noting similarities and differences. Listed also are the measurement goals and objectives (Section 5.1) aided by experimental observations from each type of monitor.

The tools discussed in this chapter are all designed for observing internal network behavior. Measurement facilities for the evaluation of distributed computing in a network environment [ALSB75, STOK76, ABRM77, GRUB77, BENN78, TREA78, USGO79] are not included. The latter set of tools are concerned only with the quality of service provided by the network, and not with probing the behavior of the network for the purpose of performance improvement.

In general, network measurement facility [MORG74] consists of:

- network traffic generator;
- software monitor;
- hardware monitor; and,
- control and data analysis software.

The different parts of a measurement facility are now discussed.

6.1 Artificial Traffic Generators

The behavior of a network is dependent on its input workload (traffic). In a measurement experiment, the traffic may come from one of two sources:

- real users; or,
- an artificial traffic generator.

With real user traffic, system behavior is observed under "real life" conditions. However, the traffic pattern cannot be controlled, and therefore the measurement experiment is not repeatable using identical input.

Artificial traffic generators, on the other hand, specify the traffic distribution exactly, and therefore the experiment is repeatable and changeable under controlled conditions. But before the measurement results can be assumed to reflect real behavior, the assumptions of the traffic distribution require validation. Studies of this nature fall into the realm of workload characterization (Chapter 7).

The artificial traffic generator must be able to simulate, under experimental control, the types of traffic the network is designed to carry. There are three different classifications of user traffic:

- interactive traffic;
- bulk traffic; and,
- stream traffic.

The traffic types are distinguished by their message size distribution, message interarrival time distribution, sensitivity to delays in message delivery, and sensitivity to data errors. A summary of the traffic types and their attributes is shown in Table 5, and explained below.

The artificial traffic generator must be able to simulate the types of traffic the network is designed to carry.

- Interactive Traffic:

  Interactive traffic [JACO77, COHE78, KITA78, TREA78] is data traffic that typically occurs in user to computer communication. A response is usually received from the host prior to the input of the next user command.

- Bulk Traffic:

  Bulk traffic [COHE78, KITA78, TREA78], such as file transfer operations, represent a flow of information at transfer speeds as fast as the network will allow.

  Bulk traffic is usually not as time critical as interactive or stream traffic, and does not always involve a user waiting for a response. An example is a file update operation in a distributed database management system.

- Stream Traffic:

  Low volume, nonbursty sources of data [JACO77, KITA78, TREA78], such as packetized voice, facsimile, and real time control commands, are examples of stream traffic. Time is the critical element: messages must be

Traffic Attributes	Traffic Types Interactive Traffic	Bulk Traffic	Stream Traffic
Message Size Distribution	short messages, large variation	long messages, large variation	constant length
Message Interarrival Time Distribution	large variation	nearly nonzero intervals	constant intervals
Sensitivity to Delays in Message Delivery	high tolerance	high tolerance	low tolerance
Sensitivity to Data Errors	low tolerance	low tolerance	high tolerance

Table 5 - Characterization of User Traffic Types

delivered within very narrow delivery constraints.

All networks analyzed handle both interactive and bulk traffic types: SATNET and KUIPNET are the only two that carry stream data as well.

Having characterized the different traffic types, it is now possible to list the basic parameters of a generator:

- message and packet size distribution;
- message interarrival time distribution;
- traffic type distribution;
- message destination;
- duration of traffic generation; and
- concurrent input emulation.

The mean and variance of the message and packet size, message interarrival time, and traffic type distribution are specified according to the desired traffic type.

In most networks, the path taken by the packet to its destination is calculated dynamically by the routing algorithm. Some networks allow user defined, fixed path routing, useful for assessing specific path characteristics.

The length of the traffic generation can be controlled either by the number of packets generated, or by a specified elapsed time interval.

Emulation can be used for two purposes: to simulate a larger network configuration, or to simulate more than one user. Each input stream operates independently in terms of packet generation and channel scheduling.

In Table 6, the artificial traffic generator parameters

Generator Parameter	A	C	D	P	S
• Message Length					
-Constant length messages	A		D		
-Random length messages	A	C			
-Single packet messages					S
-Multipacket messages					S
• Frequency of Generation					
-Constant time intervals	A		D		S
-Geometric time intervals					S
-Random (Poisson) time intervals			D		
-Generation of messages as fast as the network will allow		C			
-Generation of messages as soon as acknowledgment of the previous message is received	A				
-Change the generation frequency at prescheduled times					S
• Duration of Experiment					
-Number of messages generated			D		
-Time interval of generation				P	
• Packet Destination	A	C	D	P	S
-Forced routing		C			
• Concurrent Streams of Input				P	S

A - ARPANET [COLE72, KLEI74, KLEI76a, TOBA78]
C - CYCLADES [EYRI77, GIEN78]
D - DATAPAC [COHE78]
P - PRNET [TOBA76, TOBA78]
S - SATNET [GERL77b]

Table 6 - Artificial Traffic Generator Parameters

implemented in five packet transmission networks are summarized.

The artificial traffic generators all have implemented variations of the message and packet size distribution, interarrival time distribution, packet destination and experiment duration parameters. PRNET and SATNET are the only two networks that have concurrent stream emulation. None of the generators specify the traffic mix option.

The location of the traffic generator in the network varies: for ARPANET, PRNET and SATNET, it is part of the node software; for CYCLADES, DATAPAC and ETHERNET [SHOC79b], it resides within the network interface machine (host).

6.2 Software Monitors

A software monitor is a special program module [SVOB77, GREE79] that is part of the measured system. The monitor collects and records statistics about the state of the system, either at predefined time intervals (time driven), or whenever certain specified events occur (event driven).

The level of detail of the information collected by the monitor is flexible and easily varied. The major disadvantage is that the overhead is usually significant: the monitor competes for use of the resources that it is measuring. The only exception to this rule was found in ETHERNET [SHOC79b] where a separate, dedicated machine was attached to the network to monitor and collect data in a nonbiasing (passive) manner. This technique was possible only due to the broadcast nature of channel transmissions.

Three types of statistics are collected by software monitors for computer network performance evaluation:

- cumulative statistics;
- snapshot statistics; and
- trace statistics.

A summary of the statistics and the data measurement types (Section 5.3) they represent is given in Table 7.

The measurement goals and objectives (Section 5.1) aided by observations using each type of statistic, together with representative implementations are detailed below.

6.2.1 Cumulative Statistics

Measurements presented as sums, frequencies, histograms or matrices, represent accumulations of a variety of events over a given period of time, and are referred to as cumulative statistics [TOBA78].

These statistics describe the local behavior of a node. When all nodal statistics are analyzed together, the global network behavior measures can be calculated.

Conclusions can then be derived about the following measurement goals and objectives (Section 5.1).

- User Behavior Characterization:

  The message statistics and traffic load distribution characterize the workload that the network is expected to process, and furthermore can be used to validate assumptions made in the network design stage .

Data Measurement Type	Type of Statistic		
	Cumulative Statistics	Snapshot Statistics	Trace Statistics
Response Time and Delay Analysis: Time measurements	yes	no	yes
Resource Utilization Analysis: Space and time measurements	yes	yes	yes
Input Load Analysis: Event count measurements	yes	no	no
Resource Capacity Analysis: Length measurements	yes	yes	no
Reliability Analysis: Error measurements	yes	yes	no

Table 7 - Software Monitor Statistics and Data Measurement Types

(Section 7.2).

- System Behavior Analysis:

  The throughput and delay together provide a measure of the quality of service delivered to the user (Chapter 8). The goal of network design is to provide users with high throughput and low delay.

- Efficient Resource Allocation:

  Statistics on the utilization of network components, namely, node processors, node buffers and channels (Chapter 8), indicate how effectively the network is operating. Bottlenecks can be identified showing the need for increased resource capacity or the design of more efficient software algorithms.

- Availability and Reliability Maximization:

  Error statistics are collected to monitor the status of network components. In the event of failure, corrective action can be initiated (Chapter 10). In addition, the recoverable errors are monitored.

- Protocol Evaluation:

  To assess the effectiveness of various protocols, some special statistics are collected to analyze their behavior (Section 9.3).

The cumulative statistics collected by software monitors in four packet transmission networks are summarized in Table 8.

All networks record data regarding local and global behavior of the network. The level of detail in the

User Traffic Characterization	A	K	P	S
• Message size distribution				
-histogram of lengths	A			
-matrix of lengths sent between host pairs		K		
-histogram of last packet lengths	A			
-total words in last packets	A			
• Number of messages from each host				
-histogram	A			
-matrix of numbers sent between host pairs		K		
• Number of control messages sent to each host				
-histogram	A			
• Message intergeneration times				
-histogram			P	
Global Network Behavior				
• Number of round trips to each host				
-histogram	A			
• Round trip times				
-total	A			
-histogram			P	
• One way delay				S
Node CPU Utilization				
• Number of times background loop is executed				
-histogram		K		
Node Buffer Utilization				
• Number of times free buffer list is empty per channel	A			
• Number of buffers taken from the free list		K		
• Number of buffers returned to the free list		K		
• Number of times a buffer is unavailable (packet discarded)		K	P	
• Buffer state frequencies: free, queued for retransmission, or reserved for packet receive			P	

Table 8 - Cumulative Statistics Collected by Software Monitors

Channel Utilization	A	K	P	S
• Packet lengths in words sent per channel	A			
• Number of words sent per channel	A			
• Number of inputs received per channel	A			
• Number of conflicts per half duplex channel		K		
• I/O overhead due to channel conflicts		K		
Multiaccess Protocol Analysis				
• Number of packets received but not intended for the station			P	
• Number of transmissions per successful packet			P	
• Number of correctly received packets from each neighboring node			P	
• Number of transmissions beyond success			P	
• Number of packets incurring transmissions beyond success			P	
• Number of transmissions per successful end to end packet			P	
Error Analysis				
• Number of channel tests sent per channel	A			
• Number of channel test responses received	A			
• Number of packets received in error			P	
• Number of unacknowledged packets			P	
• Number of end-to-end unacknowledged packets			P	
• Number of duplicate packets detected			P	
• Number of alternately routed packets			P	

A - ARPANET [COLE72, KLEI74, KLEI76a, TOBA78]
K - KUIPNET [KITA78]
P - PRNET [TOBA76, KAHN78, TOBA78]
S - SATNET [GERL77b, WEIS78]

Table 8 (cont'd) - Cumulative Statistics Collected by Software Monitors statistics collected varies, and depends on the intended purpose of experiment.

6.2.2 Snapshot Statistics

Snapshot statistics [TOBA78] give the state of a node at a particular instant in time with regard to queue lengths and buffer allocation. State changes are timestamped and collected as snapshot statistics.

A summary of snapshot statistics collected by software monitors in two packet transmission networks is shown in Table 9.

The statistics collected are used to analyze buffer utilization (Section 7.2.1), and in evaluating the performance of the dynamic routing algorithms used in ARPANET [KLEI76a, KLEI77, KLEI78b] and PRNET [TOBA76, TOBA78].

6.2.3 Trace Statistics

Trace statistics are used to trace the route of packets [KLEI74, KLEI76a, TOBA76, TOBA78] as they flow through the network and measure the delays encountered at each hop.

A summary of trace statistics collected is shown in Table 10.

In ARPANET, a packet may be marked (trace bit on) for tracing as it passes through a sequence of nodes. If the trace parameter in the node is also set, a trace block is generated for each marked packet passing through.

Buffer Utilization	A	P
• Length of free storage list	A	
• Number of buffers in use for message reassembly	A	
• Number of buffers allocated for message reassembly but not yet in use	A	
Protocol Evaluation		
• Routing tables of where to send packets destined for site i	A	
• Delay table consisting of the minimum number of hops to site i plus the delay estimate	A	
• Network connectivity table		P
• Packet radio repeater states		P
• Packet radio repeater parameter values		P

A - ARPANET [COLE72, KLEI74, KLEI76a, TOBA78]
P - PRNET [TOBA76, KAHN78, TOBA78]

Table 9 - Snapshot Statistics Collected by Software Monitors

Packet Identifying Information	A	K	P
• Packet length	A		
• Source address	A	K	
• Destination address	A	K	
• Link number	A		
• Message number	A		
• Packet number	A		
• Priority/nonpriority status	A		
• Packet header, containing control information	A	K	P
Delay Measurements			
• Time the packet buffer was allocated (arrival time)	A	K	P
• Time the packet was put on an output queue	A	K	P
• Time the packet started transmission	A	K	P
• Time the packet buffer was released to the free pool (successful transmission)	A		P

A - ARPANET [COLE72, KLEI74, KLEI76a, TOBA78]
K - KUIPNET [KITA78]
P - PRNET [TOBA76, KAHN78, TOBA78]

Table 10 - Trace Statistics Collected by Software Monitors

Because of limited storage available at the packet radio units in PRNET, a new type of packet, called the "pickup packet" was devised. The pickup packets are generated with null text fields. As they flow through the network, trace statistics are stored in the text field by repeaters.

Trace data is used primarily for estimating the components of network delay (Chapter 8):

- nodal processing delay, equal to the time the packet was placed on the output queue less its arrival time;
- queueing delay, equal to the start of transmission time minus the time the packet was placed on the output queue; and,
- transmission and propagation delay, equal to the time the buffer is released to the buffer pool less the time the packet started transmission.

The trace packets are also useful for isolating looping conditions caused by the dynamic routing algorithm [MCQU78].

6.3 Hardware Monitors

A hardware monitor [SVOB77, GREE79] is a device consisting of sensors (probes) plus some processing intelligence (often a minicomputer). The probes are attached to the electronics circuitry of the measured system: signals are sensed and changes are recorded by the minicomputer.

The hardware monitor is external to the measured system, and thus does not compete for measured resources (no overhead, as in the case of software monitors). However, flexibility of the data collection is lost: the probes need to be "rewired" to change the variables recorded. Additionally, some software conditions, such as queue lengths, are difficult to detect by a hardware monitor.

Hardware monitors are usually used in parallel with a software monitor: such systems are termed "hybrid". Synchronization pulses are used [ULME78] to correlate data recorded by the hardware monitor with that of the software monitor.

In the CYCLADES network, an oscilloscope [GIEN78] was used to measure various components of propagation delay.

An IBM 2715 System Measurement Instrument (SMI) was used [ULME78] in the IBM Experimental Satellite Network to monitor the supervisor and problem state percentages of the processor, and the busy and I/O transfer states of the channel.

6.4 Network Control

Control in a network is necessary for two functions:

- experimental control, for performance assessment of the network; and,
- network monitoring, for reliability maintenance.

The techniques used to achieve each function are summarized below.

6.4.1 Experimental Control

Control of a network measurement experiment can be broken down into three tasks:

- Control Task:

  The specification of protocol and traffic parameters, enabling and disabling of statistics collection functions, and the initiation, execution and termination of the experiment are all handled by the control task.

- Collection Task:

  The collection task gathers and stores the measurement statistics for later data reduction and analysis. Two choices can be made for collecting cumulative statistics data: at the end of a time interval, a measurement packet can be sent by the node to the collecting site (host or node); or, an "examine packet" can be sent to collect the information from the node. Either method is suitable: the former is preferred because fixed time interval data is easier to analyze statistically.

- Reduction Task:

  Usually executed offline, the reduction task sorts and merges the data collected, and produces summary statistics. Reduction is best performed on the raw data so that backtracking is possible in the event of unusual or unexpected circumstances. Recording of raw data by the monitor also means less overhead is required by the measurement task.

The location of the control and data analysis software varies. In the IBM Experimental Satellite Link Network [ULME78], KUIPNET [KITA78] and PRNET [TOBA76, TOBA78], experimental control is placed in the nodes and data is logged directly onto disk or magnetic tape.

In the ETHERNET [SHOC79b], a control program finds an idle machine, loads the traffic generator, initiates execution, collected statistics are transferred to the controlling computer at the end of the experiment.

For ARPANET [KLEI74, KLEI76a, TOBA78] and the initial set of SATNET experiments [GERL77b, KLEI78, TOBA78] the Network Measurement Center (NMC) at UCLA was used. With Version 3 of the SATNET software, a new monitoring and control capability [BRES79] with an interactive user interface and improved broadcast channel measurements was installed at the site of Bolt Beranek and Newman Inc., in Cambridge, Massachusetts.

6.4.2 Network Monitoring

Status information [WEDB74, GRAN77, DURT78, NAFF78, TOBA78] is required by network operations personnel to maintain high service availability. In addition, that data is used for planning purposes, and for customer billing.

A summary of the current configuration and operational status statistics collected by existing networks is given in Table 11.

	A	C	D	P	U
Network Status					
• Up/down status of each node			D		
• Up/down status of each host	A			P	
• Up/down status of each channel	A	C	D		
• Number of absent responses to channel test signals	A				
• Number of transmission errors			D		U
• Clock values		C			
• Control panel switch settings		C			
Load Indicators					
• Number of acknowledged packets per channel	A				U
• Traffic levels of all lines			D		U
• Queue lengths		C			U
Accounting Information					
• Number of packets entering the node from each host	A				
• Number of calls to network (sessions)			D		U
• Duration of sessions			D		
• Volume of data transmitted per session			D		

A - ARPANET [MCKE72, KLEI74, KLEI76a, TOBA78]
C - CYCLADES [GRAN77]
D - DATAPAC [CLIP76, MCGI78]
P - PRNET [KAHN78, KUNZ78, TOBA78]
U - EURONET [DURT78]

Table 11 - Statistics Collected for Network Monitoring

In all networks, each node sends a "status packet" at fixed time intervals to a centralized site in the network called the Network Control Center (NCC). The NCC has the function of collecting the (possibly conflicting) reports, and determining, using heuristics, the most likely actual state of the network. In the absence of a status packet, network component failure is assumed and an alarm is posted to initiate repair activities ranging from node program reloading to contacting suppliers to correct hardware problems.

Reliability of the network components is commonly measured by the failure statistics (Section 10.2): mean time between failures (MTBF) and mean time to repair (MTTR).

Load indicators are used for long range planning, and may be used to dynamically update routing and configuration tables, or aid in flow control.

6.5 Summary, Conclusions and Open Problems

A network measurement facility generally consists of a traffic generator, software monitor, hardware monitor and control and data analysis software.

The artificial traffic generator is used to simulate user input in a controlled and repeatable manner. Generator parameters specify the characteristics of the different traffic types simulated: all network systems have implemented some variation of an artificial traffic source.

A software monitor is a program module that is part of the measured system. The most comprehensive software monitor for performance assessment was developed for ARPANET. Cumulative statistics, which gather data over a given interval of time, are the most versatile in terms of the number of different types of performance measures that can be collected. Trace statistics and snapshot statistics are specialized functions for analyzing buffer utilization, delay and routing algorithm behavior.

Hardware monitors, placed externally to the measured system, are often used in conjunction with software monitors. None of the analyzed papers reported hardware measured data.

The tasks of initiating and terminating a measurement activity, and data reduction are usually centralized at one host site. Monitoring statistics provide information about the network status, and network usage for accounting and future capacity planning purposes.

7. Network Input Process Characterization

Since a computer network [TREA78] is a communication medium, network behavior depends on traffic passing through the network. Traffic may be from real users, or artificially generated.

Computer network system models are all based on assumptions on the nature of the input process. As artificial traffic sources reflect those conjectures, real user data needs to be characterized to validate the model suppositions.

The next section presents a data stream model and identifies the relevant statistics that characterize network workload. Measurement results from real user network data traffic follow. Some open problems are discussed in the concluding section.

7.1 Data Stream Model

In this section, a prototype for the exchange of data between a user and a remote computer process, communicating via a packet transmission network, is developed. The purpose of the model is to identify the random variables that characterize a network input process.

The prototype is an adaptation of a data stream model [FUCH70, DUDI71] advanced for user terminal to computer interaction over a half duplex dedicated line. A similar, but higher level model, known as the

Stimulus-Acknowledgment-Response Model [ABRM77], was proposed to characterize a user's demands and utilization of a network to access a remote computer facility.

Modifications given below allow the prototype to be applied to packet transmission networks, and it is appropriate for process to process communication as well.

The specification of the data flow process is analogous to the early studies of Erlang [FUCH70, DUDI71] which characterized the behavior of data traffic in telephone systems. The discovered approximations of arrival rate (by the Poisson distribution) and interarrival time (exponential distribution) have retained their validity through to modern day telephone circuit switching technology, and are relevant to defining the properties of packetized voice in stream data.

Interactive traffic and bulk traffic, however, are different in that they are often characterized as "bursty". Burstiness is defined as having high peak to average ratios of demand for the channel capacity [KLEI76, LAM77, KLEI78b, LAM78], meaning that typical users make relatively infrequent requests for communication resources, but insist on considerable transmission bandwidth and switching power when needed.

A new measure, known as the bursty factor, has been defined [LAM77, EPHR78, LAM78] as the ratio of the average message interarrival time to the average message delivery delay constraint (maximum delay tolerated in message delivery specified by the user). The bursty factor forms an a priori upper bound on the channel utilization, and on the peak to average ratio.

The data exchange model describes the relationships of the flow of messages between computers and users during online or distributed processing transactions. The flow is defined in terms of random variables, but makes no assumptions on the behavior of the computational tasks themselves. Characterization of the executing procedures within a computer system is beyond the scope of the model.

First, the prototype for a user to computer session is considered. Then, the aggregation of traffic sources as occurring in packet transmission networks is introduced.

The prototype components of a half duplex single user to computer session are illustrated in Figure 3. Full duplex operation is modelled similarly, however, the time intervals may overlap since transmissions occur in both directions simultaneously.

A session, or call, is defined as the time a user initiates a data exchange process, until the end of the information exchange.

A user burst segment is the period during which a user is sending messages to the computer. It is the interval between the last character received at the terminal until the last character sent to the host (without receiving intervening messages). The first period of inactivity in a user burst segment is called user think time. Several

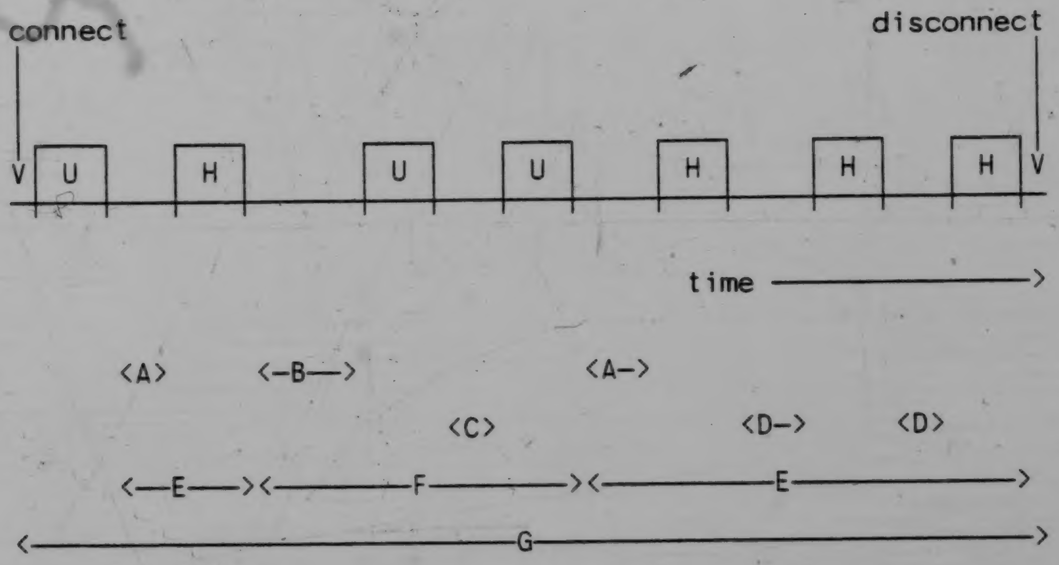
connect
disconnect
V
U
H
U
U
H
H
H
V
time
<A>
<-B->
<A->
<C>
<D->
<D>
<-E->
<-F->
<-E->
<-G->

U - user packet

H - computer packet

A - computer idle time

B - user think time

C - user interpacket time

D - computer interpacket time

E - computer burst segment

F - user burst segment

G - call or session

Figure 3 - User to Computer Data Stream Model

messages may be sent in a user burst segment: the interval between messages is the user interpacket time.

Analogously, the computer burst segment is defined with idle time (think time) and computer interpacket times.

For the single user to computer session, the statistical variables and distributions that characterize the process are:

- session time distribution;
- user burst segment length distribution;
- computer burst segment length distribution;
- number of burst segments per call;
- user packet size distribution;
- number of packets per user burst segment;
- computer packet size distribution;
- number of packets per computer burst segment;
- think time distribution;
- user interpacket time distribution;
- idle time (response time) distribution; and
- computer interpacket time distribution.

For the analysis of packet transmission networks, the individual user to computer data streams are added together to model the network input process. The data streams are assumed independent and identically distributed.

The statistics that characterize the network input process are derived by summing the random variables of the individual data streams as follows.

- Message and Packet Size Distribution:

  The user and computer packet size distributions of all data streams are summed together.

- Message Interarrival Time Distribution:

  The think time, idle time, user interpacket time and computer interpacket time distributions are totaled over all data streams.

- Session Time Distribution:

  Defined as for the data stream model, but summed over all sources.

- Number of Packets per Session:

  Amalgamated from the number of packets per burst segment for users and computers for all data streams.

In most network user traffic experiments, only the message and packet size distribution and the message interarrival time distribution are measured for performance analysis purposes. The session time and numbers of packets transferred are collected for network usage accounting systems.

As the messages exchanged between a user and host are interspersed among other messages of the input process, the user and computer burst segments are no longer meaningful.

Accumulation of statistics at the finer level of each individual user/process is often inappropriate due to the immense overhead resulting from large volumes of traffic passing through the network.

7.2 Traffic Measurement Experiments

Studies on user data traffic are summarized below. The statistics of interest, the method of data presentation and measured results are stated.

7.2.1 Message and Packet Size Distribution

Table 12 summarizes the results of four measurement experiments to investigate the message and packet size distribution of real user traffic. The observed values are listed as reported in the source documentation. For the purpose of comparison, the observation in a standard unit of measure is also supplied.

The ARPANET experiment [KLEI74, KLEI76b] measured user data traffic for a continuous seven day period during August 1-7, 1973. The network configuration consisted of about 40 nodes and approximately 50 hosts. Cumulative statistics were collected at the nodes at seven minute intervals.

The KUIPNET observation period was from April, 1976, to December, 1977. In CYCLADES, the message lengths of users [GIEN78] of one host were measured. The experimental time interval for ETHERNET was not given.

For interactive data traffic, the packet lengths for all networks was very short (about 255 bits). In ETHERNET, a bimodal distribution was observed, accounting for two types of network traffic. Most packets were short (about 240 bits), but most of the total volume of traffic was contained in file transfer (bulk) traffic (about 4400 bits). The

Data Presentation	A	C	E	K
• Number of packets per message				
-histogram (log 2 scale)	A			
-mean plotted on an hourly basis	A			
• Packet length				
-histogram (log 2 scale)	A			
-histogram of packets received (log 10 scale)		C		
-histogram of packets sent (log 10 scale)		C		
-histogram			E	K
-matrix of average lengths sent between host pairs				K
• Volume of traffic				
-histogram of packets of a given length			E	K
-histogram of number of messages				K

Measured Results

A - ARPANET [KLEI74, KLEI76a, KLEI76b]
mean packet length: 243 bits
mean of 1.12 packets per message (1008 bits per packet, 1129 bits per message)

C - CYCLADES [GIEN78]
mean packet length received: 23 octets (184 bits)
mean packet length sent: 50 octets (400 bits)

E - ETHERNET [SHOC79b]
mean packet length: 122 bytes (976 bits)
median packet length: 32 bytes (256 bits)

K - KUIPNET [KITA78]
ranges from 32 bits for predominantly interactive traffic to 1560 bits for speech traffic

Table 12 - Message and Packet Size Measurements

KUIPNET traffic matrix also showed a wide variation in the average message size, depending on the predominant type of data transferred between host pairs. Speech and picture data were almost 100 times greater in quantity and length than interactive data.

In ARPANET, the mean number of packets per message peaked during the night time hours: file transfer traffic dominated interactive traffic.

ARPANET buffers are allocated in groups of 8 to accommodate the protocol maximum of 8 packets per message. Because of a low frequency of multipacket messages, the measured buffer efficiency was .184 for the experimental period. It was shown analytically [KLEI74, KLEI76a] that reducing the buffers to 1/4 of their size could achieve a maximum efficiency of .366 for the given level of ARPANET protocol overhead. The change was not made as other performance factors would be affected.

7.2.2 Message Interarrival Time Distribution

A summary of the method of data presentation and measured results for evaluating the message interarrival time distribution is given in Table 13.

The results are from the same set of experiments described in the previous section. The ARPANET figures represent the arrival rate of messages, rather than the interarrival time distribution: the statistical relationship is that the arrival rate is the inverse of the mean

Data Presentation	A	E
• Histogram of interpacket arrival times		E
• Plot of hourly average number of host messages generated per second	A	

Measured Results

A - ARPANET [KLEI74, KLEI76a]
daytime rates were greater than nighttime rates
weekday peaks from 70 to 110 messages per second (interarrival time: 14.3 to 9.09 msec)
weekend peaks from 30 to 70 messages per second (interarrival time: 33.3 to 14.3 msec)

E - ETHERNET [SHOC79b]
heaviest load during daytime hours
mean packet interarrival time: 39.5 msec (arrival rate: 25.3 messages per second)
standard deviation: 55.0 msec
median: 8.5 msec

Table 13 - Message Interarrival Time Measurements

interarrival time.

The heaviest amount of traffic occurred during weekday daytime hours. This result is not unexpected as it represents peak activity during regular business hours. The variation is large: user think time and computer response time depend on the complexity of the task executed.

7.2.3 Path Length Distribution

Only one set of experiments investigated the path length distribution of messages.

Several statistics can be calculated directly from the network topology. For ARPANET [KLEI74, KLEI76a], they were:

- the average number of hops between node pairs was 5.32; and,
- a histogram of the number of site pairs at a given distance showed that sites were no more than 11 hops apart, and that there were more sites 7 hops apart than any other distance.

From the cumulative statistics collected, the measured results were as follows:

- the mean path length was 3.31 hops;
- the lower bound for the mean path length was calculated to be 3.24 hops, based on the assumption that all traffic flowed on the shortest paths; and,
- a histogram of the percentage of traffic travelling over a path of given length showed that 22% of the messages travelled 0 hops, 16% went 1 hop.

The measured average path length was just slightly higher than the lower bound of the mean, indicating that most of the traffic followed the shortest path.

However, the measured average path length (3.31 hops) was significantly lower than the average path length (5.32 hops) calculated from the network topology. This figure was affected by the large percentage of traffic travelling only 0 or 1 hop. The effect of "incestuous" traffic travelling between two hosts connected to the same IMP is discussed further in the next section on traffic distribution.

7.2.4 Traffic Distribution

The methods of data presentation and measured results from experiments to investigate the distribution of traffic in a network are presented in Table 14.

Traffic models used in designing communication networks often assume a uniform distribution of traffic throughout the network.

Experimental results, however, from three networks indicate that traffic is, in fact, nonuniform. This was shown by the high degree of traffic flowing between hosts connected to the same node, and by favoritism among users for a small number of hosts. The ARPANET results investigated most thoroughly the "incestuous" and "favoritism" behavior.

One hypothesis put forth [KLEI74, KLEI76a] to explain user favoritism was that once a user becomes familiar with a

Data Presentation	A	E	K
• Histogram of the normalized number of messages sent between host pairs at a given distance	A		
• Plot of the percentage of traffic flowing between hosts connected to the same node during each hour	A		
• Plot of the cumulative percentage of messages sent from the N busiest sources	A		
• Plot of the cumulative percentage of traffic between site pairs	A		
• Plot of the percentage of traffic summed over all sources to the N most favored destinations	A		
• Plot of the percentage of traffic to the most favored destination of all sources hourly	A		
• Plot of the number of favored destinations necessary to sum to 90% of the overall traffic on an hourly basis	A		
• Utilization of the network by each host		E	
• Percentage of intranetwork and internetwork traffic		E	
• Matrix of average number of messages sent between host pairs			K

Measured Results

A - ARPANET [KLEI74, KLEI76a]

38% of traffic travelled 0 or 1 hops, traffic travelling 2 to 9 hops was more uniformly distributed

an average of 22% of traffic travelled 0 hops, a peak of 80% was observed

80% of the total traffic was generated by the busiest 1/3 of the sites

90% of the traffic occurred between 12.6% of all site pairs

an average of 44% of all traffic went to favored sites

90% of the total traffic went to the 9 most favored sites

an hourly maximum of 7 favored destinations was required to achieve 90% of the overall traffic

Table 14 - Traffic Distribution Measurements

E - ETHERNET [SHOC79b]
- utilization of the network by a host varied from < 1% to over 25%
- 72% of the traffic was intranetwork traffic, 28% was internetwork traffic

K - KUIPNET [KITA78]
- 78% of the total traffic occurred between 2 hosts

Table 14 (cont'd) - Traffic Distribution Measurements particular host, he is more likely to continue using that host in the future rather than accessing others.

The ETHERNET for which measurement results are reported [SHOC79b] is joined by a high speed line to another local ETHERNET. Approximately 1/4 of the measured ETHERNET's traffic was destined for hosts attached to the other network.

When three (or more) networks are linked together, internetwork packets may travel to their destination network via another. The network(s) between the source and destination networks is called the transit network(s); the packets passing through transit networks are called transit packets.

There are no studies that measured the volume of transit packet traffic handled by a network. One study [SHOC79a], however, did explore the throughput and delay performance of a transit network using artificially generated traffic and a limited amount of real user traffic (Section 8.3).

7.3 Summary, Conclusions and Open Problems

The statistics and random variables defined for the data stream model are too detailed to be collected by a measurement facility: the overhead and volume of data would be excessive. Instead, the aggregated statistics of the network input process are collected.

The user workload (input process) characteristics are consistent across all networks measured: ARPANET, CYCLADES, ETHERNET and KUIPNET. The observed message and packet size distributions were bimodal. The peaks at approximately 255 bits and 4400 bits corresponded to average lengths of interactive and bulk traffic types, respectively. Thus, simply comparing the average packet lengths is meaningless.

The expected mean packet and message length is a parameter of network buffer management design. For ARPANET, the observed value (1.12 packets per message) was much lower then the maximum designed (8 packets per message), and thus poor buffer utilization was observed.

An assumption of uniform traffic flow throughout the network is usually made during the design and modelling phases. Measurements showed that the data flow is, in fact, not uniform.

The measured average path length (3.31 hops) was significantly lower than the mean path length (5.32 hops) calculated directly from the network topology. This was caused by a high percentage (38%) of data travelling only 0 or 1 hop.

The measured average path length (3.31 hops) was also just slightly higher than the lower bound on the mean path length (3.24 hops). Thus, most packets followed the shortest path to their destination.

No attempts were made to fit well known distribution curves to the measured data, as for example Erlang [FUCH70, DUDI71] has done. The measured data could perhaps be approximated by a mix of distributions: goodness of fit tests exist to infer the validity of such hypotheses.

In most cases, the measured results are presented in the form of mean values. Other and higher order statistics (i.e., median, mode, quantiles, variance, skewness, kurtosis, etc.) are useful for accurately specifying the distribution of a random variable.

8. Network Throughput and Delay

Throughput and delay are the two traditional indicators of performance in the evaluation of computer models and systems. In a networking environment, these system behavior variables correspond to the rate at which useful data can be transferred and to the network response time, respectively. Throughput efficiency and delay performance are measured as functions of the input load distribution.

In designing a network, a primary objective is to maximize throughput. There are three main factors [MCQU77] that affect throughput performance:

- processing capability of the node;
- amount of buffer storage space available at the node; and,
- channel capacity.

The large peak to average demand ratios (Section 7.1) that are typical in network input processes, result in low throughput averages. The excess system capacity, however, is required to meet the fast response time (low delay) demands of peak activity.

Throughput is expressed in units of packets per second (pps) or bits per second (bps). Throughput efficiency (utilization) is calculated [SUNS76, ULME78] as the ratio of useful data transmitted to the data capacity of the physical links, and is usually expressed as a percentage.

Two definitions of useful data exist: the first regards useful data as the message text only; the second includes control overhead bits incurred by the communication protocol.

Delay is defined as the interval from the start time of message transmission until the time the message arrives successfully at its destination.

The components of delay typically included in delay models [MCQU77] are defined as follows:

- Propagation Delay:

  The propagation delay is the length of time for the first bit of the packet to traverse the channel. The propagation delay is directly related to the channel capacity.

- Transmission Delay:

  The transmission delay is defined as the time for the packet to be clocked out onto the channel, and is also dependent on channel capacity.

- Nodal Processing Delay:

  The length of time it takes the node to process a packet is known as the nodal processing delay, and is dependent on the processing capability of the node.

- Queueing Delay:

  The queueing delay is defined as the length of time a packet must wait on an output queue prior to transmission. The queueing delay includes delays incurred by flow control mechanisms.

This chapter examines measured throughput and delay characteristics of packet switched and multiaccess networks.

One internetwork throughput study is also reviewed.

8.1 Packet Switched Network Performance

Tables 15 and 16 summarize the methods of data presentation and measured results of throughput and delay experiments for four packet switched networks.

In packet switched networks, the delay is approximated as it is not easily measured directly. The major obstacle is clock synchronization of the source and destination nodes. This problem is resolved by using round trip delay, i.e., the start time of message transmission until the time the acknowledgment is received for that message at the originating node. Delay is then considered to be half the round trip delay. However, this figure is only an approximation as the data and acknowledgment packets are of different lengths, and encounter different delays in traversing the network.

Among the observed networks, there is considerable variation in the measured round trip delay: this is a reflection of the differences in size and complexity of the networks studied.

KUIPNET, which had the lowest measured delay, is a network consisting of only one switching computer. For this reason, also, it was possible to measure one way delay rather than resorting to round trip delay.

ARPANET, CYCLADES and DATAPAC are large networks with many nodes widely distributed geographically. Even so, it is

Data Presentation	A	C
• Throughput as a function of the number of active traffic sources		
-message length equal to 1 packet (1008 bits)	A	
-message length equal to 8 packets (8063 bits)	A	
• Throughput as a function of message length		
* -between two nodes using 6 different packet lengths		C
* -for 5 different loop back routes using 7 different packet lengths		C
• Network average hourly channel utilization		
-including overhead	A	
-excluding overhead	A	
• Hourly channel utilization		
-of the most heavily used channel, including and excluding overhead	A	
-of the channel with the highest hourly average, including and excluding overhead	A	
-of the channel with the highest weekly average, including and excluding overhead	A	

* - measured using artificially generated traffic

Measured Results

A - ARPANET [COLE72, KLEI74, KLEI76a, KLEI76b]
- percentage of the line capacity utilized was 6.73%
- average channel utilization: 7.1% with overhead, .77% without overhead
- maximum channel utilization: 13.4% with overhead, 2.9% without overhead
- busiest line utilization: 48% with overhead, 22.5% without overhead

C - CYCLADES [EYRI77, GELE77, GIEN78]
- node throughput curves intersect when the packet lengths are the same

Table 15 - Throughput Measurements in Packet Switched Networks

Data Presentation	A	C	D	K
• Average round trip delay				
-histogram of percentages for one host (log 10 scale)		C		
• Average hourly round trip delay				
-measured using real traffic	A			
-calculated from the delay model using measured load values	A			
-calculated from the delay model using zero load values	A			
• Round trip delay as a function of packet length				
-one user session between two nodes	A			
* -for 7 different path lengths		C		
• Estimated probability density function of round trip delay				
* -for remote job entry (RJE) and time sharing (TSS) traffic under light, medium and heavy load			D	
• Average one way delay				
-histogram of number of messages				K

* - measured using artificially generated traffic

Measured Results

A - ARPANET [COLE72, KLEI74, KLEI76a]
measured average round trip delay was 93 msec

C - CYCLADES [EYRI77, GELE77, GIEN78]
round trip delay was less than 2 sec (2000 msec)
delay increased linearly with the number of nodes traversed

D - DATAPAC [COHE78]
the minimum average round trip delay was 539 msec
the delay increased slightly with link utilization

K - KUIPNET [KITA78]
average one way delay was 4 msec

Table 16 - Delay Measurements in Packet Switched Networks

difficult to compare their measured results as the networks differ greatly in protocol complexity (CYCLADES simplest, DATAPAC most complex), experimental channel speeds (DATAPAC, 2.4 Kbps, CYCLADES, 19.2 Kbps, ARPANET, 50 Kbps), and test configuration. The channel transmission speed was a dominant term in all delay models developed.

The measured round trip delay (93 msec) [KLEI74, KLEI76a] of ARPANET was found to be less than half the original design goal of 200 msec. The delay model calculations using zero input load resulted in an average hourly delay plot similar in shape to the measured average hourly delay. It was concluded that little congestion was introduced by the network itself. Using measured load values, the delay model yielded an average hourly delay plot the same shape as the measured graph, but significantly underestimated the true values. The difference was attributed to several delay factors not included in the model.

The earliest ARPANET experiments [COLE72] demonstrated the effective increase in throughput resulting from the use of an alternate path between two nodes to alleviate congestion on the primary path.

The ARPANET average channel utilization (<10%) and the utilization of the busiest channel (<50%) are low, indicating sufficient excess channel capacity for peak traffic situations.

The CYCLADES studies [GELE77, GIEN78] with varying packet lengths transmitted between two nodes showed that the throughput of each node was the same when the packet lengths were the same. When the packet lengths differed, the node transmitting longer packets experienced near optimum throughput, whereas the node sending shorter packets suffered degradation proportional to the difference in packet size. The degradation was due to channel bandwidth not utilized while waiting for the long packets to complete transmission so that acknowledgments could be sent.

The maximum throughput achieved in CYCLADES (Section 4.1.2) with the MV8 protocol was greater than with the TMM protocol using the same line speed. This was not unexpected as the MV8 protocol uses parallel logical channels to improve channel utilization. In heavy traffic, the delay increased linearly with the number of nodes traversed in the message delivery path, as expected.

The DATAPAC experiments [COHE78] consisted of a series of tests run between two minicomputer hosts connected to the same network node. A delay model was proposed, and using simulated time sharing (interactive) traffic, values for the individual delay components were measured: host processing delay, transmission delay, acknowledgment delay, nodal processing delay and window rotational delay (flow control mechanism). The major term of the delay model was found to be the transmission delay.

Through the set of DATAPAC tests simulating remote job entry (bulk) traffic, it was shown that the average delay is dependent on the frequency of collisions, i.e., acknowledgments delayed due to message transmissions in progress. The messages arrive at uniform intervals, and therefore the collisions occur at constant intervals. This in turn affects the system response time.

The DATAPAC delay model, based on a single server queueing model (M/M/1, see [KLEI76a]), is relatively simple compared to the ARPANET delay model, which is a complex equation involving many terms. The increased complexity, however, resulted in a more accurate fit (match) to the measured delay.

8.2 Multiaccess Network Performance

Tables 17 and 18 summarize the methods of data presentation and measured results for two multiaccess networks: ETHERNET and SATNET.

The definition of useful data included overhead bits in the throughput measurements. Due to the broadcast nature of the channel, it was possible to measure delay precisely (as one way delay).

As the majority of multiaccess networks are still in the experimental stages, most of the performance data, particularly for the SATNET protocols, has been measured using artificially generated traffic. For this reason, the theoretical maximum throughput, derived analytically, is included for the purposes of comparison.

Data Presentation	E	FT	RT	SA
• Channel utilization				
* -as a function of total generated load	E			
-average utilization	E			
-maximum utilization	E			
• Channel utilization as a function of the number of hosts operating at overload (100% input rate)				
* -with 5 different packet sizes	E			
• Throughput as a function of input rate				
* -3 stations, balanced traffic		FT		
* -3 stations, unbalanced traffic			RT	
* -3 stations, unbalanced traffic, symmetric gates				SA

* - measured using artificially generated traffic

Measured Results

E - ETHERNET [METC76, SHOC79b]
average channel utilization: .84%
maximum channel utilization over 1 minute: 17%
maximum channel utilization over 1 second: 37%
throughput increases linearly with load
at 98% load, throughput remains constant up to 150% load
theoretical maximum throughput efficiency: 37%

FT - SATNET Fixed-TDMA protocol [GERL77b, KLEI78a]
theoretical maximum throughput efficiency: 100%

RT - SATNET Reservation-TDMA protocol [GERL77b, KLEI78a]
utilization increases with input load: available bandwidth is assigned as requested
theoretical maximum throughput efficiency: (1-x)%, where x is the proportion of slots dedicated to routing information

SA - SATNET Slotted ALOHA protocol [GERL77b, KLEI78a]
theoretical maximum throughput efficiency: 37%

Table 17 - Throughput Measurements in Multiaccess Networks

Data Presentation	E	FT	RT	SA	CP
• Delay as a function of throughput					
* -2 stations, balanced traffic		FT	RT		CP
* -2 stations, balanced traffic, two delay classes of packets					CP
* -2 stations, unbalanced traffic			RT		
* -2 stations, balanced traffic, symmetric gates				SA	
* -2 stations, balanced traffic, asymmetric gates				SA	
* -2 stations, unbalanced traffic, symmetric gates				SA	
* -2 stations, unbalanced traffic, asymmetric gates				SA	
• Delay as a function of input rate					
* -3 stations, balanced traffic		FT			
* -3 stations, balanced traffic, symmetric gates				SA	
* -3 stations, unbalanced traffic			RT		
• Histogram of delay at selected load points					
* -2 station, balanced traffic, two delay classes of packets					CP
• Frequency of collisions	E				

* - measured using artificially generated traffic

Measured Results

E - ETHERNET [SHOC79b]

percentage of packets sent with zero queueing delay: 99.18%

percentage of packets delayed due to carrier sensing: .79%

frequency of collisions: .03%

FT - SATNET Fixed-TDMA protocol [GERL77b, KLEI78a]

RT - SATNET Reservation-TDMA protocol [GERL77b, KLEI78a]

SA - SATNET Slotted-ALOHA protocol [GERL77b, KLEI78a]

CP - SATNET CPODA protocol [CHU78]

near perfect agreement with theoretical and simulation results were observed for all SATNET protocols, see Table 17 for more details

Table 18 - Delay Measurements in Multiaccess Networks

Measured over a 24 hour period with real traffic, the ETHERNET average channel utilization was low (<1%). The maximum utilization observed over 1 second intervals was 37%. Using artifically generated traffic, the throughput increased linearly with load until the maximum utilization was achieved. The utilization then remained constant up to 150% load.

To evaluate the delay characteristics of ETHERNET, a test program was written to periodically transmit a packet. The frequency of packet collisions was recorded. With the channel utilization so low (<1%), it is not surprising that over 99% of the packets were sent with zero queueing delay, and only .03% were involved in collisions.

For all SATNET protocols (Section 4.2.4), near perfect agreement with theoretical and simulation results were observed for balanced traffic situations.

The F-TDMA experiments were performed primarily to calibrate all the SATNET measurements. The heavy load delay model analysis, based on Little's result, is sufficiently general [GERL77b, KLEI78a] as to be applicable to any protocol scheme operating at saturation. This was used as a reference point for validating the operation of other protocols.

The R-TDMA unbalanced traffic experiment demonstrated the ability of stations to acquire unused slots. As the input rate increased, the available bandwidth was progressively allocated.

The S-ALOHA studies showed that the optimal gate values depend on the traffic pattern. This is the motivation for the controlled S-ALOHA protocol which adjusts the gate values dynamically according to the input load.

To adequately compare the delay and throughput characteristics of protocols, an exhaustive study of protocol behavior under different traffic pattern combinations is required. So far, one traffic case has been done [GERL77b]. The F-TDMA, R-TDMA and S-ALOHA protocols were evaluated using a two station configuration, with one node heavily loaded, the other varying the input traffic rate.

For that case, R-TDMA yielded the highest throughput and lowest delays for the input traffic range of .01 to .4 packets/slot. For input rates lower than .01 packets/slot, S-ALOHA was best, but only marginally better than R-TDMA. The F-TDMA protocol was preferred for input rates > .4 packets/slot, with R-TDMA as the second choice.

In the CPODA protocol, reservations may be piggybacked on message headers, thus decreasing total delay slightly as system load increases. However, the degree of improved delay performance reduced as the number of stations increases. This was due to a lower probability of message header reservation piggybacking.

CPODA also allows delay class specification for messages. For the longer delay class, the reponse time remained constant as the load increased. The shorter delay class consumed the entire delay advantage of reservation piggybacking. The variance of the shorter delay class response time was smaller than that of the larger delay class. Therefore, the shorter delay class packets are more likely to be delivered within their delay class constraints.

8.3 Internetwork Performance

To conclude the discussion of throughput and delay, the preliminary results of an internetwork study are presented.

The case study [SHOC79a] involved the use of RADIONET as a transit network between two ETHERNET implementations. The ETHERNETs were also joined together by a high speed leased line.

In the absence of international standards for internetworking (levels 1-4, Section 3.3), a gateway computer performed the necessary interface between the two networks. The approach was to encapsulate, i.e., treat as data, the ETHERNET protocol within the RADIONET protocol.

Differences in packet structure and protocol function required special attention in the network interface. First, the host addressing scheme in the three networks was standardized to facilitate the RADIONET routing algorithm. Secondly, the ETHERNET protocol allows broadcast addressing, where one packet has multiple destination addresses. This function was simulated for messages destined for RADIONET hosts by generating multiple copies of the packet, one for each destination. Finally, packet length differences led to fragmentation: one maximum size ETHERNET packet fragmented into three RADIONET packets. Fragment sequencing and reassembly was a responsibility of the gateway.

The internetwork throughput experiments were conducted in December, 1978, and January, 1979. Using an artificially generated traffic stream, the throughput was measured for seven configurations (locations) of the source and destination processes.

The performance was "disappointing" [SHOC79a] in that the throughput efficiency of the RADIONET link was only 12% compared to the 86% efficiency of the direct leased line link. The overhead for handling each packet was high, indicating that the internetwork system was limited by the processing power of the packet radio unit (PRU).

8.4 Summary, Conclusions and Open Problems

The fundamental performance measures of throughput and delay are used to compare network systems and models.

Throughput measurements are presented in terms of how much data (including overhead) can be transmitted per unit of time. Delay is measured directly as elapsed time in multiaccess networks, and estimated by round trip delay in packet switched networks.

As networks are designed to meet the dual objectives of high throughput efficiency and low delay, both parameters must be examined using tradeoff relationships.

Throughput measurements based on real user input have been presented for only two networks: ARPANET and ETHERNET. A network design objective is to provide sufficient channel capacity to meet the response time demands of peak activity: the average throughput efficiency is therefore expected to be low. For both ARPANET and ETHERNET, the observed channel utilizations (throughput efficiencies) were low (7.1% and .84%, respectively), because of light traffic in addition to the planned excess capacity.

Packet switched networks are the best understood in terms of behavior and performance. Multiaccess networks are not widely used as yet, so few results with real user traffic are available.

The throughputs of ETHERNET and SATNET were measured using artificially generated traffic to simulate high load. The observed maximum throughput efficiencies closely matched the theoretical maxima derived from analytic models (Table 17). For the S-ALOHA and CSMA protocols, it is only possible to achieve 37% throughput efficiency due to packet collisions. In the F-TDMA protocol, however, collisions do not occur, and thus 100% throughput efficiency is attainable.

Delay measurements for real user traffic were given for five networks. The KUIPNET and ETHERNET local area networks both had fast response times (4 msec and <1 msec, respectively). The wide area networks of ARPANET, CYCLADES and DATAPAC all had response times of less than 2 seconds, acceptable for most network users. The CYCLADES response time (<2000 msec), however, was much greater than ARPANET (93 msec) or DATAPAC (539 msec) even though the geographic area serviced is much smaller. Since differences in response time are due mainly to protocol complexity, channel speeds and test configurations, direct comparison of delay is very difficult.

Internetwork studies are just beginning: problems of protocol incompatibility need to be resolved by the adoption of international standards.

Any assessment must consider, in addition to throughput and delay, a much broader spectrum of requirements. The specification of the input workload is crucial, as the relative performance is dependent on the traffic characteristics, such as the mix of interactive, bulk and stream traffic types. Additionally, the level of service provided affects the evaluation as the complexity can vary from a simple transport service (such as CYCLADES) to a fully automated, transparent and reliable system (such as ARPANET or DATAPAC, or any X.25 compatible network), and hence should be considered as well.

9. Fairness, Stability and Protocol Efficiency

In the previous chapter, throughput was defined as the rate at which useful data is transferred through the network. The definition of useful data included overhead bits.

The throughput of multiaccess networks was investigated from a total system behavior point of view. However, due to the nature of the channel access scheme, the effects of variations in the input load distribution on throughput are more pronounced. Therefore, the throughput of each node is also studied using composite measures known as fairness and stability. Fairness verifies that the network resources are allocated proportionally according to each node's demand; stability refers to the rate at which the system returns to average (normal) throughput after a period of peak activity.

Also in this chapter, the nature of the protocol overhead is investigated, and discussed in terms of its effects on throughput. In other words, the difference between the two types of useful data is quantified.

9.1 Fairness

A protocol may be efficient in terms of throughput and delay, and yet be unfair [JACO78, KLEI78a, TOBA78] in allocating channel bandwidth according to each node's demand.

To assess fairness of a protocol, the ratio between the input rate and the throughput of each station is examined to determine if the following criteria are satisfied:

- all backlogged stations get an equal share of the channel;
- all stations not backlogged get a share of the channel equal to their input traffic requirements;
- the throughput of nonbacklogged stations cannot exceed the throughput of backlogged stations; and,
- the sum of the individual channel allocations is equal to the maximum network capacity.

In packet switched networks, the channels (transmission lines) are shared exclusively between each node pair, and thus are not subject to problems with fairness. In multiaccess networks, on the other hand, many stations access the channel, and therefore analysis of fairness is particularly important for these networks.

Table 19 summarizes the measurement experiments conducted to analyze fairness in ETHERNET (Section 4.2.2) and three protocols in SATNET (Section 4.2.4). As the F-TDMA protocol is guaranteed fair by its equal assignment of bandwidth to all stations, no test results were recorded.

The R-TDMA studies revealed that the protocol is fair: unused bandwidth is progressively acquired until the channel is equally apportioned.

As long as balanced traffic conditions prevailed and symmetric gate values were used, the S-ALOHA protocol demonstrated fairness. However, when the gate values were

Data Presentation	E	RT	SA	CP
• Throughput of each host				
-10 hosts, each offering 10% load (high load)	E			
-90 hosts, each offering 100% load (overload)	E			
• Throughput of each stations as a function of input load				
-2 stations, balanced traffic		RT		
-2 stations, unbalanced traffic		RT		
-3 stations, unbalanced traffic		RT		CP
-3 stations, balanced traffic, symmetric gates			SA	
• Delay of each station as a function of input load				
-2 stations, balanced traffic		RT		
-2 stations, unbalanced traffic		RT		
-3 stations, unbalanced traffic		RT		CP
-3 stations, balanced traffic, symmetric gates			SA	
• Delay as a function of throughput				
-2 stations, balanced traffic, symmetric gates			SA	
-2 stations, balanced traffic, asymmetric gates			SA	

Note: Artificially generated traffic was used in these studies

Measured Results

E - ETHERNET [SHOC79b]

high load test: average throughput efficiency was 94%, individual host utilizations varied from 9.3% to 9.6%

overload test: average host utilization was 1.1%, the range was from .9% to 1.3%

RT - SATNET Reservation-TDMA protocol [GERL77b, KLEI78a]

delay performance was identical for all stations

slots were equally assigned among stations

unused bandwidth was progressively assigned as the input rate increased until the channel was equally apportioned

Table 19 - Fairness Measurements

SA - SATNET Slotted-ALOHA protocol [GERL77b, KLEI78a]
- protocol is fair if symmetric gates and balanced traffic conditions prevailed
- asymmetric gates and unbalanced traffic resulted in capture situations
- controlled S-ALOHA tests prevented the capture situations, i.e., the protocol was fair

CP - SATNET CPODA protocol [CHU78]
- because S-ALOHA was used for making channel reservations, capture effects were demonstrated

Table 19 (cont'd) - Fairness Measurements asymmetric, particularly when PN << PR, then a capture situation arose. One station retained almost exclusive control over the channel (capture situation) for an extended period of time, where the choice of station was random.

The tests were rerun using the controlled S-ALOHA algorithm in which the gate values are adjusted dynamically according to load. For these runs, the results demonstrated fairness of the protocol.

The definition of fairness was altered slightly for the CPODA protocol to require that the channel capacity be shared equally among users in the same delay class and priority. The S-ALOHA discipline is used to make channel reservations. The experiment did not show the expected capture situation, but a further study using simulation was able to demonstrate the capture effects. The use of a control algorithm would allow fair usage of the channel by all stations.

9.2 System Stability

During unfavorable (peak) traffic conditions, network systems are driven into a degraded mode of operation [KLEI78a, TOBA78]. Stability is assessed by the speed at which the system recovers to equilibrium throughput and delay: persistence of low throughput and long delay indicate unstable behavior.

The performance measures that are used to analyze stability are:

- average throughput prior to the pulse (peak load);
- time to recover after the pulse; and,
- average throughput during the recovery period.

The results of measurement experiments to evaluate system stability are summarized in Table 20.

The ETHERNET system was shown stable as throughput did not decline with increases in the input load. Experimental results suggest that proposed control schemes [TOKO77] offer increased complexity with little potential benefit in throughput recovery.

The S-ALOHA protocol throughput behaved as predicted analytically, requiring 20,000 slots (10 minutes) to recover from a pulse 20 slots (600 msec) long. Dynamic control of the gate values by the controlled S-ALOHA protocol reduced the recovery period to 1000 slots (30 seconds). Thus the controlled S-ALOHA protocol is much more stable than the uncontrolled S-ALOHA protocol.

9.3 Protocol Efficiency

Protocol efficiency is the ratio of useful (message text only) traffic to the total traffic carried by the network [KLEI76b, SUNS76]. Protocol overhead is that fraction of the total traffic that is not user message text data. Overhead depends on the protocol implementation, and consists of all the necessary bits and packets required to transport a message between processes.

Data Presentation	E	SA
• Channel utilization as a function of total offered load	E	
• Channel utilization as a function of the number of hosts operating at overload -using 5 different packet lengths	E	
• Throughput as a function of time -10 stations, symmetric gates slots 0-1300: 2 stations operating at 100% load slots 1301-1320: 10 stations operating at 100% load (pulse) slots > 1320: 2 stations operating at 100% load		SA

Note: Artificially generated traffic was used in these studies

Measured Results

E - ETHERNET [SHOC79b]
throughput (98%) did not decrease with load increases

SA - SATNET Slotted-ALOHA protocol [KLEI78a, TOBA78]
throughput during the pulse degraded from .5 packets/slot to .01 packets/slot
to clear the backlog, 20,000 slots (10 minutes) were required
using controlled S-ALOHA system recovery was attained after 1,000 slots (30 seconds)

Table 20 - Stability Measurements

Overhead can be divided into four distinct, additive categories [KLEI76b]:

- node to node control overhead;
- source node to destination node control overhead;
- host to host control overhead; and,
- background control overhead.

The first three types of overhead are required to monitor each packet as it flows through the network. The amount of overhead per packet is constant, thus line (protocol) efficiency increases with packet length. The total amount of overhead due to each of these three types will increase proportionally with the total volume of user data traffic.

The fourth type of overhead, background control, consists of network status and routing information. The total amount of type four overhead in the network remains constant, regardless of the user input load distribution.

In Table 21, protocol efficiency in terms of overhead measurements for ARPANET and ETHERNET, based on real user traffic flow, are summarized.

The ARPANET observations [KLEI76b], measured over a one week period in May, 1974, are given as two figures: the first is the actual observed value of overhead. Due to the low volume of user traffic, channel utilization was very low (6.73%). The second figure given in Table 21 represents the projected overhead at 100% line utilization.

The ratio of overhead bits to each user message bit is

Data Presentation	A	E
• Average percentage of overhead		
-node to node control overhead	A	
-source node to destination node control overhead	A	
-host to host control overhead	A	
-background control overhead	A	
-all types of overhead		E
• Average percentage of user data	A	E
• Frequency distribution of network control messages (node to node plus source node to destination node control overhead)	A	
• Frequency distribution of host to host control overhead messages	A	

Measured Results

A - ARPANET [KLEI76a, KLEI76b]

node to node control overhead: 13.2% of the total traffic (35.22% projected)

source node to destination node control overhead: 9.1% of the total traffic (24.43% projected)

host to host control overhead: 4.7% of the total traffic (12.59% projected)

background control overhead: 64.2% of the total traffic (4.32% projected)

user data (nonoverhead bits): 8.8% of the total traffic (23.44% projected)

percentage of line capacity utilized: 6.73% (100% projected)

most common network control message: "RFNM-S" (88.77%) acknowledging receipt of a single packet message

most common host to host control message: "ALL" (63.8-79.0%) allocating buffers at the receiving host

E - ETHERNET [SHOC79b]

user data (nonoverhead bits): 69%

overhead bits: 31%

relative to the average packet length (122 bytes):
  protocol encapsulation: 4 bytes (3% overhead)
  internetwork protocol: 22 bytes (18% overhead)

relative to the median packet length (32 bytes):
  protocol encapsulation: 4 bytes (12.5% overhead)
  internetwork protocol: 22 bytes (69% overhead)

Table 21 - Protocol Overhead Measurements

3.1 to 1. At a finer level of detail, the node to node overhead was highest (1.5 to 1), source node to destination node overhead second (1 to 1), and host to host overhead lowest (.5 to 1). This proportionality was retained in the projected figures.

In the observed traffic measurements, the background overhead bits dominated user bits (7.3 to 1). Since the volume of background bits remains constant, the ratio reduces significantly (.2 to 1) at projected high load.

Thus, while the observed ARPANET protocol efficiency was only 8.8%, the projected upper bound is 23.44%.

The relative frequencies control messages (network and host to host) were given. The network control message sent most often (88.77%) was "RFNM-S", signalling correct receipt of a single packet message. This high proportion of single packet messages agrees with message and packet size distribution results discussed earlier (Section 7.2.1). The most common host to host control message sent (63.8-79.0%) was "ALL", signalling the allocation of buffer space for a message at the receiving host.

The ETHERNET results [SHOC79b] are significantly different from those observed for ARPANET. The ratio of overhead bits per user data bit was .45 to 1 (ARPANET: 3.1 to 1). The relative simplicity of ETHERNET has produced an efficient protocol. However, the ARPANET complexities of flow control and routing allow multichannel traversal of messages. The relative efficiency of multiaccess and packet switched protocols, as observed, was predicted in [KLEI78b].

The clue to the difference in the level of overhead observed to that given for an average length packet can be found in the bimodal distribution of packet sizes (Section 7.2.1). Since the median packet length is much shorter than the average packet length, there are more small packets than long packets. The small packets have higher overhead, and therefore reduce protocol efficiency.

9.4 Summary, Conclusions and Open Problems

In multiaccess networks, input variations have a significant effect on throughput and delay, thus requiring assessment of protocol fairness and stability as well. A protocol is considered "fair" if network resources are allocated proportionally accoring to each node's demand, and "stable" if the recovery period to average (normal) throughput after a peak load condition is short.

The S-ALOHA protocol was found to be both unfair and unstable as capture situations and prolonged throughput degradations were observed. The controlled S-ALOHA protocol, with its dynamic gate value adjustment algorithm, resulted in a fair and stable protocol; the expense of this increased protocol complexity was increased processing overhead at the node.

The CPODA protocol uses the S-ALOHA discipline to make reservations during the reservation subframe. The measurement data did not show the expected unfairness and instability. However, a simulation study demonstrated the capture effects, suggesting that controlled S-ALOHA should be used instead of S-ALOHA.

The ETHERNET protocol (CSMA) was found both stable and fair.

Protocol efficiency examines the proportion of the total user traffic that is message text data only. Overhead is defined to include all bits and packets required by the protocol to manage the transfer of data.

Total overhead in the ARPANET protocol consists of a constant amount of overhead due to background packets, plus the overhead associated with each message packet. ETHERNET overhead consists of only message packet overhead.

The ARPANET and ETHERNET protocol efficiencies were observed to be 8.8% (23.44% projected) and 69%, respectively. Because the background overhead in ARPANET remains constant regardless of load, the protocol efficiency at 100% load was calculated. The background overhead efficiency dropped from 64.2% to 4.32%. Of the message packet overhead, the node to node component was greatest (13.2% observed, 35.22% projected).

The efficiency of a protocol decreases as more control and service functions are provided. The ETHERNET protocol is relatively simple, and therefore highly efficient. The ARPANET protocol, providing functions of flow control, routing and multichannel traversal of messages, is considerably more complex and thus requires an order of magnitude more overhead to transfer user messages.

With the increased levels of overhead defined for distributed processing (Section 3.3) in an Open System Interconnection (OSI) environment, protocol efficiency can be expected to decrease even further. OSI designers must be careful to minimize the overhead associated with process to process communication over a network.

10. Network Reliability

This chapter concentrates on network reliability. Reliability issues [MCQU77] fall into two broad categories: reliability of service, and reliability of network connectivity (also termed availability [GRUB77]).

Service reliability performance is defined as the network's ability to successfully, and accurately, deliver the message. The capability of the network to detect and correct errors is tested.

In the next section, the sensitivity of protocols to channel noise is analyzed by examining average (bit and message) error rates and effective throughput. Other measures of service reliability are:

- percentage of messages delivered;
- detected error rate; and,
- undetected error rate.

In all cases, the ability of the protocols to detect and correct errors is tested.

Reliability of network connectivity is analyzed to evaluate the probability of hardware component failures. The objective is to maximize the mean time between failure (MTBF) and minimize the mean time to repair (MTTR).

Other measures of availability are:

- the number of nodes and/or lines necessary to disconnect the network; and,
- the fraction of node pairs not connected by any path.

Availability of a network can be improved significantly using hardware redundancy design techniques. A discussion of these methods is beyond the scope of this thesis.

10.1 Sensitivity to Channel Noise

The sensitivity of a protocol to noise on a channel is crucial to efficiency and reliability of the communication network. This is particularly true for protocols that rely on past channel observations and measurements for the scheduling of future transmissions, or for updating parameters.

Studies of service reliability are summarized in Table 22. ARPANET, DATAPAC and ETHERNET report measured results for cable-based channels with real user traffic. The SATNET satellite channel studies utilized artificially generated traffic.

The ARPANET channels [KLEI74, KLEI76a] showed a wide variation in performance, from no errors to an error rate of 1 in 340 packets.

The DATAPAC observations [COHE78] determined that the message arrival rate does not affect the channel error rate. Additionally, successful error recovery was achieved at the frame level (OSI level 2); the packet level (OSI level 3) was not impacted at all.

The ETHERNET error rate [SHOC79b] was so low that packets received in error were excluded from other reported performance results.

Data Presentation	A	D	E	FT	RT	SA	CP
• Average error rate	A	D	E				
• Number of channels having a given packet error rate	A						
• Throughput as a function of noise level							
* -2 station, balanced traffic							CP
* -2 station, unbalanced traffic				FT	RT		
* -3 station, unbalanced traffic					RT		
* -20 station (emulated), 50% noise level						SA	
• Delay as a function of noise level							
* -2 station, balanced traffic							CP
* -2 station, unbalanced traffic					RT		
* -3 station, unbalanced traffic					RT		

* - measured using artificially generated traffic

Measured Results

A - ARPANET [KLEI74, KLEI76a]
average packet error rate: 1 in 12,880 packets (7.8×10^{-5} errors/packet)
more than half the channels had an error rate $< 10^{-5}$
14% of all channels had no errors
the noisiest channel error rate: 1 in 340 packets (2.9×10^{-3} errors/packet)

D - DATAPAC [COHE78]
average error rate: 1 in 400,000 bits (packet size was 101 octets, 2.02×10^{-3} errors/packet)

E - ETHERNET [SHOC79b]
average error rate: 1 in 2,000,000 packets ($.5 \times 10^{-6}$ errors/packet)

FT - SATNET Fixed-TDMA protocol [GERL77b]
throughput degradation is negligible for noise gate values up to 3%

RT - SATNET Reservation-TDMA protocol [GERL77b, KLEI78a]
at 30% noise level, R-TDMA throughput was still superior to F-TDMA

Table 22 - Reliability Measurements

SA - SATNET Controlled S-ALOHA protocol [KLEI78b]
throughput with 50% noise levels using the dynamic gate adjustment algorithm was only 10% lower than optimum

CP - SATNET CPODA protocol [CHU78]
delay increases and throughput drops significantly when the bit error rate is > 10^{-5}

Table 22 (cont'd) - Reliability Measurements

A comparison of the measured results suggests that the channel error rate is most likely proportional to the length of cable. This would account for the two orders of magnitude difference in error rates of ETHERNET and ARPANET/DATAPAC.

The SATNET satellite channel experiments [GERL77b, CHU78, KLEI78a] all used artificial noise gates to evaluate the level of sensitivity. Noise gates are parametric values in the station arbitrarily set at initialization time to indicate the fraction of received packets to be discarded (simulated errors). Routing packets are assumed immune to noise.

Unbalanced traffic conditions are used to analyze protocol robustness to noise. Load imbalance emphasizes the performance degradation resulting from loss of synchronization and packet collisions.

Degradation in the performance of the F-TDMA protocol due to noise will result only from packet retransmissions. The scheduling algorithm does not depend on observations of network behavior.

The R-TDMA, S-ALOHA and CPODA protocols are all distributed scheduling algorithms that depend on global network behavior information to make their scheduling decisions. They are thus expected to be sensitive to errors.

Noise causes R-TDMA to revert to the F-TDMA scheduling discipline if synchronization with other stations is lost. Therefore, F-TDMA was expected to outperform R-TDMA beyond some critical noise gate value. At 30% noise levels, the performance of R-TDMA was still superior to F-TDMA.

The controlled S-ALOHA protocol was shown analytically insensitive to noise. Throughput with 50% noise levels was only 10% less than optimum.

The CPODA protocol is extremely sensitive to noise as acknowledgments and reservations may be placed in the message packet headers. Each CPODA station, therefore, is required to correctly hear all transmissions to properly schedule packets. Throughput and delay deteriorated rapidly when the bit error rate was > 10^{-5}.

10.2 Availability

The observed availability statistics for network hardware components are given in Table 23. The ARPANET measurements are based on a 19 month interval; the CYCLADES data on over 2 years of network operation.

The MTBF figures for the nodes and channels indicate that the hardware components are fairly reliable, with availability > 95%.

The CYCLADES channel failure distribution [GRAN77] suggests a high correlation: some lines failed regularly, while others not at all. The same result most likely holds for ARPANET as well, judging from the variation in channel error rates observed in the previous section.

Statistics for MTTR are not given as they are dependent on the nature of the maintenance contract with the hardware vendor.

Data Presentation	A	C
• Mean time between failures (MTBF)		
-node (cpu)	A	C
-channel	A	C
• Histogram of the number of nodes having a given percentage of down time	A	

Measured Results

A - ARPANET [KLEI74, KLEI76a, GRUB77]
- the average node down time was 1.64%, with a worst case behavior of 9.13%
- the MTBF for nodes was 403 hours
- the average circuit down time was .38%

C - CYCLADES [GRAN77]
- the node MTBF was 8 months, failures were usually attributed to hardware problems
- the channel MTBF was about 1 year
- the channel failure distribution was highly correlated, with some lines failing regularly and others none at all

Table 23 - Availability Measurements

10.3 Summary, Conclusions and Open Problems

In analyzing service reliability of a network, it is the protocol error detection and recovery procedures (such as packet sequencing and numbering, or bit error detection and correction codes) that are tested. For those protocols, such as CPODA, that depend on observations of network behavior to schedule future transmissions, the degree of sensitivity to channel noise affects the level of throughput degradation.

The observed average channel error rates for three cable based networks, ARPANET, DATAPAC and ETHERNET were 7.8×10^{-5}, 2.02×10^{-3} and $.5 \times 10^{-6}$ errors per packet, respectively. The error rate is proportional to cable length, accounting for the two orders of magnitude difference between the ETHERNET local area network, and the long haul networks of ARPANET and DATAPAC. It is interesting to note that the maximum channel error rate for ARPANET (2.9×10^{-3}) is about the same as the average channel error rate for DATAPAC.

The SATNET satellite channel experiments employed artificially generated traffic together with simulated noise levels. The F-TDMA protocol does not depend on network observations for scheduling, and therefore is immune to noise. The R-TDMA, controlled S-ALOHA and CPODA protocols, however, are expected to be noise sensitive. Experimental results showed that R-TDMA outperformed F-TDMA even at 30% noise levels, and with controlled S-ALOHA the throughput efficiency was degraded by only 10% at 50% noise levels. CPODA, which is most dependent on system observations for scheduling, was the most sensitive to noise, degrading significantly at bit error rates > 10^{-5}.

Reliability of network connectivity (availability) is measured by the frequency of hardware component failure. The failure statistics of mean time between failure (MTBF) and mean time to repair (MTTR) are used. The relevant measurement papers contained very little availability data.

11. Summary, Conclusions and Future Research

The objective of this thesis was to examine, compare and contrast measurement techniques used in the performance evaluation of computer communication networks.

A subset of a comprehensive bibliography on computer network design and performance analysis (Chapter 2), containing references to publications reporting measurement data, formed the source material for the analysis.

Definitions, the types of networks, international protocol standardization activities (Chapter 3), and an overview of some specific networks and their protocol implementations (Chapter 4) established the reference base for the remainder of the thesis.

The goals and objectives of a network measurement experiment (Chapter 5) were identified. The hardware and/or software tools needed for the monitoring facility (Chapter 6) were presented in terms of the types of statistics collected, and the measurement objectives satisfied in different networks. An artificial traffic generator is used to simulate user data traffic types under controlled and repeatable conditions. A software monitor collecting cumulative statistics is the most useful performance measurement tool.

A model for the network input process was developed, and observed values of the workload variables from different networks were compared (Chapter 7). Packet and message sizes had bimodal distributions, hence comparisons of mean lengths are meaningless.

The network performance measures were then considered by comparing results from various systems using real or artificially generated traffic.

Throughput and delay (Chapter 8) are the basic evaluation parameters. High throughput and low delay are dual network performance objectives, and thus tradeoff relationships must be examined. In multiaccess networks, the protocol fairness and stability (Chapter 9) also need to be studied.

Protocol efficiency (Chapter 9) depends on the relative complexity of the protocol and network services provided. For example, ETHERNET was more efficient than ARPANET, but ARPANET provides multihop message transfers.

The sensitivity of a protocol to noise (Chapter 10) is studied in terms of throughput degradation. The CPODA protocol was found the most sensitive to noise.

In planning a network facility, the Open System Interconnection (OSI) model is recommended for wide area networks. For local area networks, the ETHERNET design is suggested.

In this chapter, the results from previous chapters are interpreted in terms of the future of computer communication. The following topics are considered:

- the role of standardization;
- the proposed new network services;
- the impact of technological advances;

- the requirements of performance measurement tools;
- the nature of the network workload; and,
- the expected network performance.

Finally, the relevance of a bibliographic database as a tool for research is assessed.

11.1 Standardization

In Section 3.3, a design, known as the Open System Interconnection (OSI), was presented as the proposed standard for process to process communication in a packet switched network environment. To date, the first three levels of the seven layer architecture have been accepted internationally as the CCITT X.25 Recommendation; the timeframe for completing the specification is December, 1980 [BACH78a].

However, to meet this optimistic schedule, considerable cooperation is required from participating members of the regulatory bodies. Charles Bachman (Chairman of the ANSI Study Group on Distributed Systems, and Chairman of the International Standards Organization Technical Committee on Open System Interconnection) alluded to the enormity of the task when he stated [BACH78a]:

> "These schedules are "impossible" by all past experience, except for the example of CCITT's X.25. It is considered that the schedule is technically feasible. The problems lie in marshalling the people and gaining the release of

> the currently proprietary information, which can serve as the technical basis for the standards. The negotiations will be difficult because there will be alternative means to solve many of the problems. Selecting one among several viable choices is the most difficult task for any committee which needs to act by concensus."

It remains to be seen how the standards for distributed systems will evolve.

11.2 New Network Services

In the area of new network services, many different activities are taking place:

- Expansion of packet switched networks:

  There are several new packet communication services at the implementation stages in different countries (Section 4.1). Most are adhering to the X.25 Recommendation for packet transport service functions.

- Expansion of multiaccess networks:

  Multiaccess networks utilizing radio and satellite channels are still in the experimental stages. Local area networks using cable as the transmission medium are more advanced. As yet there are no standards for channel access methods, and so are an issue for international concensus. Additionally, an open area for research is network design using more than one channel for multihop transmissions.

- Internetworking:

  Interconnection of X.25 compatible packet switched networks is readily achieved using the principles of the OSI.

  To date, only one effort has been recorded [SHOC79b] in joining together multiaccess networks. Interconnection takes the form of protocol encapsulation (the source protocol is treated as data in the transit network). Standardization efforts would reduce incompatibilities between different protocols.

  The linkage of packet switched and multiaccess networks again takes the form of protocol encapsulation. Results from ARPANET-SATNET experiments can be expected in the near future.

Users can expect global access to computer systems as soon as regulatory and political issues can be resolved.

11.3 Technological Advances

Technological upgrades to higher capacity channels and faster processing speeds for nodes are the trend for network components. The technologies on the verge of widespread usage [JACO78, KAHN78, RAWS78, ROBE78, PARK79] are:

- Satellite channels:

  Higher density satellite channel transmissions and multibeam broadcasting are expected enhancements. The introduction of onboard processing power will enable the satellite to act as a node, and, eventually control intersatellite transmissions. The advent of space shuttle technology provides for the repair of satellites, thus improving reliability performance.

- Radio channels:

  The primary motivation for developing packet radio networks is to provide a mechanism for digital land mobile communication. The viability of radio networks depends on technological improvements in radio receivers and transmitters, and miniaturization of components to reduce bulk (to pocket calculator size). The concepts of multihop broadcasting still need investigation.

- Optic fibers:

  Optic fibers are the technological replacement for wire cables, having significantly greater capacity per unit cost, and comparatively smaller size and weight.

- Large scale integration:

  Large scale integration technology has resulted in higher packing densities in electronics circuitry. Dramatic improvements in cpu processing speeds have been achieved, a definite asset to network nodes. As communication standards evolve, hardware integration of protocol functions can also be expected.

Significant shifts in cost/performance ratios could lead to a re-examination of network topology and node processor design. The ultimate goal is to service more users more quickly and more effectively.

11.4 Network Measurement Tools

The range of measurement tools for existing packet networks is extensive. A wide variety of statistics are defined for characterizing virtually every aspect of network behavior. However, not all monitors utilize the same units of measure, nor do they all implement the full set of observation functions. The ARPANET network is the most comprehensive in its measurement capabilities.

With the emergence of process to process communication protocols, measurement tools will need to be redesigned to match more closely the functional layers of the OSI [NAFF78]. Inefficiencies of each level can then be isolated and analyzed.

The level of overhead incurred by a measurement facility must be minimized to avoid biases in the data. This concern will be even more relevant as the complexity of protocols increase.

11.5 Network User Workload

The packet size and message interarrival time characteristics of interactive user data traffic were found to be fairly uniform among networks (Section 7.2).

However, the proportion of interactive traffic to other types of traffic will change in the future. As distributed systems become common, the volume of file transfer (bulk) data will increase. Transaction-oriented traffic, from automated inventory and electronic funds transfer applications, will introduce many sessions of short duration with small packets. Digital voice traffic, teleconferencing, and real time computer applications (stream traffic) will also increase.

From the point of view of network performance, the characterization of the input process will become more complex. Networks will introduce concepts such as message priority and broadcast addressing to meet the different performance requirements of the different traffic types. Evaluations of effectiveness of the new features (algorithms) will be of interest to both network designers and users.

11.6 Network Performance

It is difficult to compare packet networks solely on measured performance values of throughput, delay, cost and reliability. Other factors must also be taken into consideration when evaluating a computer communication network service:

- Transparency:

  Transparency is the ability of a network service to deliver the data exactly as received. Code conversion at the network protocol level removes illegal bit sequences, such as those which may cause accidental disconnects. Higher level protocols (for example, levels 5-7 of OSI, Section 3.3) are concerned with data representation differences (such as computer word length) between heterogeneous computer systems.

- Privacy and security:

  The network must be able to prevent unauthorized modification or disclosure of information it has been entrusted to transfer.

- User friendliness:

  User friendliness is measured by the length of time required to learn a system, and the quality of diagnostic aids available for online/offline debugging of problems.

- Interface requirements:

  The level of effort required by the user or host to achieve a network connection may range from a simple telephone call to a local node, to the purchase of specialized hardware, to the development of software to format data according to protocol specifications.

Technological innovations and protocol sophistication do not alter the fundamental definitions of throughput, delay, cost and reliability. However, the relative absolute values will most certainly change, not only with respect to a network's implementation versions, but also relative to other networks. Hardware upgrades, for example, should ideally reduce delay and improve throughput, with greater reliability at less cost.

11.7 Epilog

Research and development in understanding the behavior of computer networks has progressed substantially since the concept of packet communication was first proposed in the mid 1960's. The intent of this thesis was to analyze the performance measurement aspects of designing, implementing and operating a network.

Political, sociological and regulatory issues are not addressed. For a comprehensive set of references on the topics of network design and performance analysis, the reader is referred to the *Computer Network Performance Bibliography* [TOWN80a]. The database provides a useful starting point for network communication research.

The future of computer communication is bright. We are on the threshold of many new and exciting developments, limited only by the imagination of the network designers, and the challenge of tomorrow.

Annotated Bibliography

[Note: Some references are included that are not explicitly referred to in the text. They are relevant to the study of computer network performance and are included for the sake of completeness.]

ABRM77 Abrams, M.D., Cotton, I.W., Ward Watkins, S., Rosenthal, R., Rippy, D.E., "The NBS Network Measurement System". IEEE Transactions on Communications, Vol. COM-25, No. 10, Oct. 1977, pp. 1189-1198.

The Network Measurement Machine (NMM) is a tool that characterizes a user's demands and utilization of a network to access remote computing services. Internal transport network behavior is not investigated: only the global measures of response time, throughput and turnaround time are collected. The emphasis of the paper is on the design of the hardware and software monitors, and the data analysis package. 12 references

ABRN73 Abramson, N., "Packet Switching with Satellites". National Computer Conference, New York, New York, June 4-8, 1973.

The historical development of data communication, covering telephone systems, ARPANET, and the radio based ALOHA system, led to the proposal of using satellite channels as a transmission medium. The properties of satellite channels are given, and the channel capacity using pure ALOHA and slotted ALOHA access schemes are derived. Consideration is given to users accessing the channel at arbitrary data rates. 20 references

ALSB75 Alsberg, P.A., "Distributed Processing on the ARPA Network -- Measurements of the Cost and Performance Tradeoffs for Numerical Tasks". Hawaii International Conference on System Sciences, Honolulu, Hawaii, January 7-9, 1975, pp. 91-94.

The cost of networking is considered in relation to a distributed processing environment. Jobs are sent to the optimal (cheapest) resource for that class of computational tasks. 2 references

BACH78a Bachman, C.W., "Domestic and International Standards Activities for Distributed Systems". IEEE Computer Society International Conference: CompCon Fall 1978, Washington, D.C., September 5-8, 1978, pp. 140-143.

The historical development of standards activities for distributed systems, leading to the creation of the International Standards Organization (ISO) subcommittee on "Open System Interconnection" is given. Three working groups are examining the "overall architecture", "transport service" and "users of transport services". The related activities of CCITT (international standards body), ANSI (United States), BSI (United Kingdom), and AFNOR (France) are put in perspective with the ISO subcommittee work. 2 references

BACH78b Bachman, C.W., Canepa, M., "The Session Control Layer of an Open System Interconnection". IEEE Computer Society International Conference: CompCon Fall 1978, Washington, D.C., September 5-8, 1978, pp. 150-156.

The functions of the session control layer to establish, maintain and terminate logical connections between processes in an Open System Interconnection are discussed. Covered also are the issues of addressing, data transfer, integrity, security and control administration. 3 references

BENN78 Bennett, C.J., Hinchley, A.J., "Measurement of the Transmission Control Protocol". Computer Networks, Vol. 2, 1978, pp. 396-408.

A set of experiments to evaluate the implementation efficiency, throughput and delay characteristics of the Transmission Control Protocol (TCP) is presented and discussed. TCP is an end to end protocol that provides acknowledgment, sequencing and flow control function for internetwork communications. The experimental tools developed measured only response times at the source and destination processes using timestamping. Measurement of the transport network behavior was not done. Throughput was found limited by channel capacity, however, elimination of duplicated functions in the TCP and transport protocols would have further improved throughput. 19 references

BIND75 Binder, R., Abramson, N., Kuo, F., Okinaka, A., Wax, D., "ALOHA Packet Broadcasting -- A Retrospect". National Computer Conference, Anaheim, California, May 19-22, 1975.

The ALOHANET packet broadcasting radio network is described. The protocol issues of random access channel control, broadcast channel control, packet length, addressing, and error control are resolved, with the rationale behind each decision given. Choices concerning the design of the radio communication subsystem are examined. The impact of microprocessor technology on the user interface is assessed, and areas for future research given. ALOHANET was the first network to successfully utilize the packet broadcasting concept for online access of a central computer via radio. 15 references

BLAC78 Blackshaw, R., Open Systems Architecture: What is it, and why is it required?, Unpublished paper, Sept., 1978.

This paper discusses the motivation and design philosophy of the Open Systems Architecture. Excellent analogies to standards activities in telephone networks and computer programming languages are given. Illustrations given in the Appendix conceptualize the operation of the layered architecture. no references

BLEV78 Blevanus, J., Mitaut, G., "ESOPE: A Tool for Qualitative Testing of TRANSPAC External Procedures". European Computing Congress: Eurocomp 78, London, England, May 9-12, 1978, pp. 775-788.

To assist in debugging TRANSPAC network subscriber (host) interface software, a tool called ESOPE was developed. ESOPE consists of a small minicomputer system, containing an artificial traffic source able to generate under parametric control "erroneous" data of all types (acknowledgment, sequence, frame format errors, etc.). An observation function analyzes network behavior under these conditions. An online monitor is used to configure the test scenarios. Extensive detail is provided on the modular design and function of each component of ESOPE. no references

BRES79 Bressler, R.D., Heart, F.E., "The Atlantic Packet Satellite Experiment". Pacific Telecommunications Conference, Honolulu, Hawaii, January 8-9, 1979, pp. 1B-45 to 1B-51.

The goals and configuration of the SATNET packet satellite broadcast network begin this paper. The three versions of the Satellite IMP (SIMP) software, with their resulting enhancements, are described. The paper ends with an overview of the new SATNET and Gateway Monitoring and Control Centers (SMCC, GMCC). 13 references

CHU78 Chu, W.W., Naylor, W.E., "Measurement and Simulation Results of C-PODA Protocol Performance". National Telecommunications Conference, Birmingham, Alabama, December 3-6, 1978, pp. 4.2.1-4.2.7.

Initial experiments with the Contention Based Priority Oriented Demand Assignment Protocol (CPODA) on the SATNET satellite broadcast network are presented. The SATNET experimental environment is described, and the CPODA protocol discipline reviewed. Performance measurements evaluated delay, throughput, fairness, and robustness to noise. The effect of frame size and minimum control subframe size on throughput and delay was investigated. Also analyzed was the ability of CPODA to schedule different delay classes. 8 references

CLAR78 Clark, D.D., Pogran, K.T., Reed, D.P., "An Introduction to Local Area Networks". Proceedings of the IEEE, Vol. 66, No. 11, Nov. 1978, pp. 1497-1517.

A tutorial paper, the concepts pertinent to local area networks are reviewed. The distinctions and similarities to long haul networks and to computer I/O bus structures are enumerated on the topics of network topology, transmission control structure, reliability characteristics, data flow, addressing scheme, transmission medium, and protocol design. The nature of the hardware and software structure of the host interface is investigated: a case study of the Local Network Interface (LNI) is presented. Finally, the interconnection of networks is discussed in terms of the subnetwork concept. 34 references

CLIP76 Clipsham, W.W., Glave, F.E., Narraway, M.L., "Datapac Network Overview". *International Conference on Computer Communication,* Toronto, Ontario, August 3-6, 1976, pp. 131-136.

The motivation for providing packet switched services is reviewed in relation to the development of the DATAPAC network. The subscriber interfaces for hosts and intelligent terminals (called Standard Network Access Protocol, or SNAP), and asynchronous terminals (called Interactive Terminal Interface, or ITI) are described, as well as the message priority, collect call, and closed user group features. The design of the SL-10 data switch, employing a multiprocessor architecture is presented, together with performance objectives for network transit delay and error statistics. Finally, the role of the Network Control Center is covered. 5 references

COHE78 Cohen, N.B., Field, J.A., Kalra, S.N., "Measurement of X.25 Service". *IEEE Computer Society International Conference: COMPCON Fall 1978,* Washington, D.C., September 5-8, 1978, pp. 330-337.

The measurement configuration at the University of Waterloo for evaluating the performance of the DATAPAC network is outlined. The emphasis of the paper is on measurement results: call setup time, error characteristics, acknowledgment behavior, and round trip delay. An artificial traffic generator is used to simulate remote job entry (RJE) and time sharing (TSS) traffic. A delay model is presented and analyzed for time sharing traffic conditions. 10 references

COLE72 Cole, G.D., "Performance Measurements on the ARPA Computer Network". *IEEE Transactions on Communications,* Vol. COM-20, No. 3, June 1972, pp. 630-636.

The network measurement tools developed for use with the ARPANET are described in detail: accumulated statistics, snapshot statistics, trace statistics, and artificial traffic generation. The tools are applied to a set of two experiments between two nodes. The objectives were to evaluate round trip delay variations according to traffic type, and to examine the effects of adaptive routing on throughput. 3 references

DESJ78 desJardins, R., White, G., "ANSI Reference Model for Distributed Systems". IEEE Computer Society International Conference: CompCon Fall 1978, Washington, D.C., September 5-8, 1978, pp. 144-149.

The ANSI reference model of Open System Interconnection is presented. The seven layer architecture is discussed with respect to functionality of each layer, and the parallel standards and/or implementations already in place. The concepts and motivations are explained. 9 references

DUDI71 Dudick, A.L., Fuchs, E., Jackson, P.E., "Data Traffic Measurements for Inquiry-Response Computer Communication Systems". IFIP Congress: Information Processing 71, 1971, pp. 634-641.

The work of [FUCH70] is extended to an analysis of the data stream model of inquiry-response systems characterized by very short interactive sessions, such as online banking and production control applications. 4 references

DURT78 Durteste, B., "Network Control and Monitoring for an International Public Network (EURONET)". European Computing Congress: Eurocomp 78, London, England, May 9-12, 1978, pp. 985-991.

The proposed architecture for EURONET and functions of each structural unit are given. The network monitoring and control requirements are listed together with the measurement statistics collected at each component site. Error detection and recovery procedures are outlined. Only a small subset of the total architecture will be implemented by EURONET initially. no references

EMMO78 Emmons, W.F., "Network Control Protocols". IEEE Computer Society International Conference: CompCon Fall 1978, Washington, D.C., September 5-8, 1978, pp. 238-245.

Summarized are the efforts and proposals of ISO and ANSI for code independent (bit oriented, rather than character oriented) headers in the transport service layers of the Open System Interconnection. The headings achieve independence and are transparent to the higher control levels (users of transport services). no references

EPHR78 Ephremides, A, "On the "Bursty Factor" as a Measure for Characterizing Data Traffic". IEEE Transactions on Communications, Vol. COM-26, No. 11, Nov. 1978, pp. 1791-1792. [Includes author's reply]

It is argued that the bursty factor presented in [LAM78] is not a new measure, but merely a transition between channel capacity determining delay to delay determining channel capacity. The author replies that the delay requirement is indeed inherent in the traffic source as poor response time will lead to discontinued use of the network. 1 reference

EYRI77 Eyries, F., Gien, M., "On Line Performance Measurement in the CYCLADES Computer Network". European Conference on Electrotechnics: EUROCON '77, Venice, Italy, May 3-7, 1977, pp. 3.1.7(1-6).

An overview of the CIGALE packet transport layer of the CYCLADES computer network is followed by a brief description of network measurement tools developed. Experimental results of maximum throughput and round trip delay as functions of packet length and route length are presented. 8 references

FRAN72 Frank, H., Kahn, R.E., Kleinrock, L., "Computer Communication Network Design -- Experience with Theory and Practice". Spring Joint Computer Conference, Atlantic City, New Jersey, May 16-18, 1972.

The ARPANET network design problems are addressed by specifying the objectives and solution techniques used in three key areas: node (IMP) design, topological design and network performance. The IMP properties are considered in terms of message handling, error control, flow control and routing. Network topology is developed with reliability requirements and optimized subject to cost constraints. Finally, analytic models based on single server and networks of queues theory, are used to predict delay, and calculate optimum channel capacity assignments. 52 references

FUCH70 Fuchs, E., Jackson, P.E., "Estimates of Distributions of Random Variables for Certain Commputer Communications Traffic Models". Communications of the ACM, Vol. 13, No. 12, Dec. 1970, pp. 752-757.

A data stream model of interactive user to computer traffic is developed. Results of a set of experiments measuring four timesharing systems are discussed, and a goodness of fit analysis identified the probability distributions that best approximate the random variables. The conclusions suggested that the model was robust, tractible, and suitable for input into more complex models. 28 references

GELE77 Gelenbe, E., Grange, J.L., Mussard, P., Performance Limits of The TMM Protocol: Modelling and Measurement. Institut de Recherche d'Informatique et d'Automatique, Rocquencourt, France, Report No. 230, April 1977.

Throughput measurements of a two node CIGALE network configuration using the TMM protocol are given. A model of the protocol behavior is validated against the measured results and found accurate. State transition diagrams are used to deduce the semi Markov process. Analytically, it was found that the maximum throughput cannot exceed 2/3 of the two way channel capacity, and that the optimum values of the relative packet sizes can be calculated. 6 references

GERL77a Gerla, M., Kleinrock, L., "Closed Loop Stability Controls for S-ALOHA Satellite Communications". Data Communications Symposium, Snowbird, Utah, September 27-29, 1977, pp. 2-10 to 2-19.

The S-ALOHA protocol is subject to stability problems. Previous work involving the Control Limit Policy or Retransmission Control Procedure was shown inadequate for networks utilizing multiple buffers or time-varying traffic patterns. A Closed Loop Control approach, based on the observation that maximum throughput is attained when the sum of packet transmission and retransmission rates is equal to one, is investigated. Three algorithm variations are compared: complete knowledge, collision detect, and collision nondetect. Simulation studies looked at the rate of convergence to the steady state solution, and the dynamic gate value adjustment behavior under pulse and cyclic traffic patterns. The performance of the algorithm does not depend on the accuracy of the collision estimate. 10 references

GERL77b Gerla, M., Nelson, L., Kleinrock, L., "Packet Satellite Multiple Access: Models and Measurements". *National Telecommunications Conference,* Los Angeles, California, December 5-7, 1977, pp. 12:2-1 to 12:2-8.

A set of experiments testing the performance of three multiaccess protocols on the SATNET satellite broadcast network are described. The objectives of the measurement experiment are given. Cumulative statistics based on artificially generated traffic were collected at the UCLA Network Measurement Center. Delay models were developed and throughput versus delay measurements were made for Fixed-TDMA (F-TDMA), Reservation-TDMA (R-TDMA), and Slotted-ALOHA (S-ALOHA) protocols. For the S-ALOHA protocol, the effect of asymmetric gate values and unbalanced traffic was investigated. The fairness and robustness to noise of R-TDMA was demonstrated. The paper ends with a brief comparison of the protocols, and directions for future research. 11 references

GIEN78 Gien, M., Grange, J., "Performance Evaluations in CYCLADES". *International Conference on Computer Communication,* Kyoto, Japan, September 26-29, 1978, pp. 23-32.

This paper summarizes results from performance measurement of the CYCLADES network. First, an overview is given of the CYCLADES architecture and the CIGALE transport service. The study of propagation delay using a hardware monitor is discussed. Throughput and delay analysis of the TMM and MV8 protocols using artificially generated traffic over a variety of node configurations and routes follow. Finally, the message length distribution and round trip delay of real user traffic from one host are given. [Includes a summary of GELE77.] 30 references

GRAN77 Grange, J., "Operating the CIGALE Packet Switching Network: Concepts, Techniques and Results". *European Conference on Electrotechnics: EUROCON*

'77, Venice, Italy, May 3-7, 1977, pp. 3.1.6(1-6).

The advantages and disadvantages of centralized versus distributed network control are examined. The CYCLADES network uses centralized control for its CIGALE packet transport service to monitor the network status. Monitoring at the node assists in adaptive routing and congestion control computations, and in the event of software/hardware failure, enables dynamic reloading of the system. Diagnostic tools developed for failure analysis are listed, as well as reliability figures based on over two years of operational data. 19 references

GREE79 Greene, W., Pooch, U., "Distributed Network Performance Monitoring: A Critical Review". *Hawaii International Conference on System Sciences,* Honolulu, Hawaii, January 4-5, 1979, pp. 259-271.

The performance monitoring of computer networks is assessed from an objective point of view. The reasons for installing monitoring systems are given, and followed by a discussion of the different types of monitors: software, hardware and hybrid. The design criteria of communication monitors are reviewed, and a methodology for implementation and testing of a hybrid system is presented. Based on material contained in [MORG74, WEDB74, SVOB77]. 23 references

GRUB77 Grubb, D.S., Cotton, I.W., "Criteria for Evaluation of Data Communications Services". *Computer Networks,* Vol. 1, 1977, pp. 325-340.

Nine criteria are identified for evaluating the quality of service provided by a telecommunication facility (all types, including packet transport networks) for the interchange of data between processes. Each characteristic is defined, and discussed in terms of its significance to the evaluation. The criteria are: transfer rate, availability, reliability, accuracy, channel establishment time, network delay, line turnaround delay, transparency and security. Typical values are provided where possible. 35 references

JACO77 Jacobs, I., Lee, L., Viterbi, A., Binder, R., Bressler, R., Hsu, N., Weissler, R., "CPODA -- A

Demand Assignment Protocol for SATNET". Data Communications Symposium, Snowbird, Utah, September 27-29, 1977, pp. 2-5 to 2-9.

The Contention Based Priority Oriented Demand Assignment (CPODA) protocol is described, together with the rationale for the design decisions. The goals and objectives of the protocol are given with respect to the user traffic and hardware environment: packetized voice and data, message priority and delay classes, and earth stations with different receiving rates. The technique for making channel reservations on a contention basis (slotted-ALOHA) is detailed. The service ordering algorithm and delay considerations are discussed. The issues of stability control, synchronization and message acknowledgment are addressed. CPODA assigns only the channel bandwidth required to transmit the message, rather than using a slotted channel structure. CPODA will be implemented on the SATNET satellite broadcast network. 16 references

JACO78 Jacobs, I.M., Binder, R., Hoversten, E.V., "General Purpose Packet Satellite Networks". Proceedings of the IEEE, Vol. 66, No. 11, Nov. 1978, pp. 1448-1467.

The concept of a General Purpose Packet Satellite Network (GPSN) is introduced in terms of its requirements. The characteristics of satellite channels, and the operation of the transponder and earth station hardware are summarized. The shortcomings of conventional demand assignment access schemes are itemized, followed by an extensive discussion on the more powerful, but more complex, design of the PODA protocol. The hardware and software configuration of the SATNET testbed is given with an overview of research activities carried out to date. Finally, the future of satellite technology is projected. 37 references

KAHN78 Kahn, R.E., Gronemeyer, S.A., Burchfiel, J., Kunzelman, R.C., "Advances in Packet Radio Technology". Proceedings of the IEEE, Vol. 66, No. 11, Nov. 1978, pp. 1468-1496.

The basic concepts of a digital radio network are presented in tutorial form. The functional requirements of a radio network are listed, and the basic design options of broadcast channel access, configuration, store and forward operation, routing and mobile operation are considered. The advanced features of dynamic reconfiguration, multihop and multidestination routing are also discussed. The PRNET project is reviewed as a case study: the network monitoring and measurement facilities are described. A major emphasis of the paper is on the hardware design of the packet radio unit and the use of spread spectrum signalling technology. 39 references

KITA78 Kitazawa, S., Sakai, T., "Performance Evaluation of the Kuipnet Computer Network". Computer Communications, Vol. 1, No. 3, June, 1978, pp. 149-155.

The local area Kyoto University information processing network (KUIPNET) configuration is presented together with the four types of traffic the network handles: real time data, picture processing data, file transfer data and interactive data. The network measurement machine and data analysis programs are described. Results of IMP performance, traffic characteristics, and host to host control command behavior observations are assessed. 15 references

KLEI73 Kleinrock, L., Lam, S.S., "Packet-Switching in a Slotted Satellite Channel". National Computer Conference, New York, New York, June 4-8, 1973.

The use of satellite channels for packet switching with the pure ALOHA, slotted ALOHA, and reservation ALOHA access schemes is discussed. An analytical investigation into the tradeoff between average delay and throughput for slotted ALOHA is presented using both uniform and nonuniform traffic sources. 13 references

KLEI74 Kleinrock, L., Naylor, W.E., "On Measured Behavior of the ARPA Network". National Computer Conference, Chicago, Illinois, May 6-10, 1974.

The results of a 7 day measurement experiment using the ARPANET network are presented and analyzed. The measurement tools, consisting of a trace facility, accumulated statistics, snapshot statistics, and status reports are used to collect data. Characterized are: message and packet length distribution, path length distribution, message interarrival time distribution, traffic type distribution, overall traffic distribution, channel utilization, packet error distribution, IMP failure rate distribution, and round trip delay. 18 references

KLEI76a Kleinrock, L., *Queueing Systems, Volume II: Computer Applications.* Toronto, Ontario: John Wiley & Sons, Inc., 1976.

The central theme of this book is the application of queueing systems to computers. The first three chapters provide the theoretical foundation, chapter four analyzes time sharing systems. The last two chapters (half the book) concentrate on networks. Chapter five considers problems of design, covering topology, capacity assignment, flow assignment and delay analysis with reference to ARPANET, satellite and ground radio packet switching. Chapter six discusses the performance measurement and analysis of ARPANET. The areas evaluated are: routing, flow control, lockups, degradation, traps, throughput, user input characteristics, and protocol overhead. [Includes COLE72, FRAN72, MCQU72, KLEI73, KLEI74, KLEI76b, KLEI77.] 142 references in Chapters 5 and 6, not all unique.

KLEI76b Kleinrock, L., Naylor, W.E., Opderbeck, H., "A Study of Line Overhead in the ARPANET". *Communications of the ACM,* Vol. 19, No. 1, Jan. 1976, pp. 3-13.

The components of the ARPANET protocol overhead are evaluated on the basis of a week of packet trace data. The overhead is subdivided into four levels of control: node to node, source node to destination node, host to host, and background. Measured results are given together with a projection of performance at peak load. An equation for line efficiency is derived, and analyzed under various allocation and data length conditions. Finally, a comparison is made between the ARPANET protocol and the TCP internetwork protocol. 23 references

KLEI77 Kleinrock, L., Opderbeck, H., "Throughput in the ARPANET - Protocols and Measurement". *IEEE Transactions on Communications,* Vol. COM-25, No. 1, Jan. 1977, pp. 95-104.

The behavior of the ARPANET multipacket message handling, in terms of throughput degradation, is the focal point of this paper. Using artificial traffic generated at maximum rates, two major protocol design problems were discovered. The first concerned "loop traps" created by the loop detection procedure of the adaptive routing algorithm. The second problem, called "phasing", resulted from conflicts in the assignment of buffer resources whereby piggybacked allocations superceded explicit buffer requests. 13 references

KLEI78a Kleinrock, L., Gerla, M., "On the Measured Performance of Packet Satellite Access Schemes". International Conference on Computer Communication, Kyoto, Japan, September 26-29, 1978, pp. 535-542.

The set of results presented by Gerla, Nelson and Kleinrock [GERL77] are extended for three and higher station configurations. As before, the experiments were to evaluate multiaccess protocols for the SATNET satellite broadcast network. Analytical delay models follow a description of the protocols. Delay and throughput measurements of Fixed-TDMA (F-TDMA) protocol calibrated the experiment. Reservation-TDMA (R-TDMA) was investigated with respect to throughput, delay, fairness and robustness to noise. Fairness and stability analyses were performed for Slotted-ALOHA (S-ALOHA). Controlled S-ALOHA, designed to solve the stability problems of S-ALOHA, was tested for the effect of the protocol parameters (window size and loop gain factor) on throughput as a function of input rate. Fairness, stability and robustness to noise were also tested. 16 references

KLEI78b Kleinrock, L., "Principles and Lessons in Packet Communications". Proceedings of the IEEE, Vol. 66, No. 11, Nov. 1978, pp. 1320-1329.

This paper is intended to be a reflection on the development of packet communication in the past decade. The problem of distributed resource allocation and sharing is defined; the reasons for the failure of conventional analysis techniques in providing packet network design solutions are given. Drawing upon examples from the ARPA sponsored projects (ARPANET, SATNET and PRNET), principles are derived for the areas of deadlock prevention, performance degradation, broadcast channel access schemes, flow control, routing control and hierarchical network configurations. 27 references

KUNZ78 Kunzelman, R.C., "Overview of the ARPA Packet Radio Experimental Network". IEEE Computer Society International Conference: COMPCON Spring 1978, San Francisco, California, February 28-March 3, 1978, pp. 157-160.

An overview of the packet radio experimental network (PRNET) is given, with the physical configurations of two implementations of the network. The protocols for hop by hop control, station to terminal control, and internetwork data transfer are outlined. The issues of message routing, packet acknowledgment and retransmission, and system monitoring are addressed. PRNET is the first radio system to combine multihop routing capabilities with multiaccess shared radio channel techniques. 11 references

LAM77 Lam, S.S., "Satellite Multiaccess Schemes for Data Traffic". International Conference on Communications, Chicago, Illinois, June 12-15, 1977, pp. 37.1-19 to 37.1-24.

The use of satellite channels for computer data traffic is assessed. A quantitative measure, called the bursty factor, is introduced for characterizing data traffic sources. Multiaccess schemes are classified into three groups: channel reservation, random access and packet reservation. Throughput and channel capacity bounds are given for representative protocols from each category. Finally, the delay performance is compared under a variety of traffic environments. 19 references

LAM78 Lam, S.S., "A New Measure for Characterizing Data Traffic". IEEE Transactions on Communications, Vol. COM-26, No. 1, Jan. 1978, pp. 137-140.

A measure called the bursty factor is introduced to characterize the burstiness of a data traffic source. The measure relies only on the traffic source characteristics and the message delivery delay constraints, independent of any specific network design. The mathematical relationships between the bursty factor, channel capacity, channel utilization and peak to average ratio for single and aggregate traffic sources are derived. 6 references

LELA76 Le Lann, G., Le Goff, H., "Advances in Performance Evaluation of Communication Protocols". International Conference on Computer Communications, Toronto, Ontario, August 3-6, 1976, pp. 361-366.

Simulation studies are used to evaluate the throughput and delay performance of several interprocess and internode protocols in the CYCLADES network. The CK and ZE interprocess protocols are examined under bulk and interactive traffic conditions. Three variations of the MV8 internode protocol are then compared. 6 references

MACD78 MacDonald, V.C., "Domestic and International Standards Activities for Data Communications". IEEE Computer Society International Conference: CompCon Fall 1978, Washington, D.C., September 5-8, 1978, pp. 253-260.

The functions and roles of two international standards bodies (CCITT and ISO) are contrasted. The mandate of each is presented with respect to a layered architecture for network communication. 10 references

MCGI78 McGibbon, C.I., Gibbs, H., Young, S.C.K., "Datapac - Initial Experiences with a Commercial Packet Network". International Conference on Computer Communication, Kyoto, Japan, September 26-29, 1978, pp. 103-108.

The development of DATAPAC from initial implementation, field trial to commercial service is presented from an operational point of view. The functions of the Network Control Center (NCC) and Data Collection Center (DCC) are outlined. Future plans are given in the areas of nodal expansion, internetworking, network interface machines, and new services. 8 references

MCKE72 McKenzie, A.A., Cosell, B.P., McQuillan, J.M., Thrope, M.J., "The Network Control Center for the ARPA Network". International Conference on Computer Communications, October 24-26, 1972, pp. 185-191.

The function and operation of the Network Control Center (NCC) in ARPANET is the topic of this paper. The network monitoring mechanisms for detecting node (IMP) and line failures, and the heuristics used to determine the most likely actual state of the network are discussed. Finally, experiences in operating the network over a period of one year are related. 6 references

MCQU72 McQuillan, J.M., Crowther, W.R., Cosell, B.P., Walden, D.C., Heart, F.E., "Improvements in the Design and Performance of the ARPA Network". *Fall Joint Computer Conference*, Anaheim, California, December 5-7, 1972.

An excellent background paper, the source to destination flow control, source to destination sequence control, and IMP to IMP transmission control procedures are defined. The organization of the IMP programs, in terms of data structures, packet flow and processing, and module functions, are documented. Finally, a brief analysis of throughput, round trip delay and line utilization is considered. 13 references

MCQU77 McQuillan, J.M., Walden, D.C., "The ARPA Network Design Decisions". *Computer Networks*, Vol. 1, 1977, pp. 243-289.

A review of the design of ARPANET, this paper begins by addressing the fundamental properties and requirements of packet communication networks. Analytic models and tradeoff relationships are derived for the performance measures of delay, throughput, cost and reliability. The choices of network equipment (node processors, channels, and host to node connections) are matched to specifications. The routing procedure and node to node transmission procedure are assessed with respect to impact on network design and performance. The designation of responsibilities between the network and host is evaluated in terms of the end to end protocol. 36 references

MCQU78 McQuillan, J.M., Falk, G., Richer, I., "A Review of the Development and Performance of the ARPANET Routing Algorithm". *IEEE Transactions on Communications*, Vol. COM-26, No. 12, Dec. 1978, pp. 1802-1811.

The paper is a survey of the development of the ARPANET distributed adaptive routing algorithm. Modifications to the original implementation are motivated by examples from operational data and explained by state diagrams. Network disturbances caused by the lack of built in congestion control led to the definition of a revised routing algorithm. 8 references

METC76 Metcalfe, R.M., Boggs, D.R., "Ethernet: Distributed Packet Switching for Local Computer Networks". *Communications of the ACM,* Vol. 19, No. 7, July 1976, pp. 395-404.

The relationship of local computer networking to remote computer networking and multiprocessing is given as an implementation between the two extremes. An excellent system summary of the ETHERNET multiaccess local network, utilizing a coaxial cable as the transmission medium, is presented. The issues of design principles, hardware and software implementation, growth potential, expected performance and interprocess communication protocol are all covered. 36 references

MORG74 Morgan, D.E., Banks, W., Colvin, W., Sutton, D., "A Performance Measurement System for Computer Networks". *IFIP Congress: Information Processing 74,* Stockholm, Sweden, August 5-10, 1974, pp. 29-33.

The objectives of network monitoring systems, and classification of the types of data collected are listed. Computer system monitors available are compared to the functional requirements of a network monitor system: none satisfied all criteria. The University of Waterloo network monitoring system design and implementation details are then given. 33 references

NAFF78 Naffah, N., "Diagnostics and Supervision in Informatics Networks". *European Computing Congress: Eurocomp 78,* London, England, May 9-12, 1978, pp. 647-658.

With the advent of distributed systems based on the Open System Interconnection (OSI) concept, this paper proposes the structuring of monitoring tasks to parallel the network architecture. The supervisory functions associated with each level are isolated, and error detection and possible recovery techniques are listed. 27 references

PARK79 Parkhill, D.F., "The Future of Computer/Communications". Infor, Vol. 17, No. 2, May 1979, pp. 187-203.

The information society of tomorrow is projected in terms of technological advances and socio-economic impact. Large scale integration, optic fibers and satellites are the trends in network hardware components. The VIEWDATA system is presented as a case study, together with scenarios of the future. 6 references

POUZ74 Pouzin, L., "CIGALE, The Packet Switching Machine of the CYCLADES Computer Network". IFIP Congress: Information Processing 74, Stockholm, Sweden, August 5-10, 1974, pp. 155-159.

The CYCLADES network is overviewed with respect to design options, and definitions of network components and procedures. The division of tasks between CIGALE, the packet transport service, and CYCLADES, the higher level host support, is rationalized. Finally, the requirements for interconnecting CYCLADES with other networks are presented. 28 references

POUZ79 Pouzin, L., "Recent Developments in Data Networks". Conference 1979, Quebec, Quebec, June 25-28, 1979, pp. 299-305.

This paper begins with a comprehensive inventory and current status report of world computer communication networks offering circuit switching or packet switching services. The issue of tariffs and impact of standardization is assessed in terms of future network developments. no references

RAWS78 Rawson, E.G., Metcalfe, R.M., "Fibernet: Multimode Optical Fibers for Local Computer Networks". IEEE Transactions on Communications, Vol. COM-26, No. 7, July 1978, pp. 983-990.

The technical difficulties of using optic fibers as the transmission medium in local area network configurations (bus, star or ring) are described. A project, known as FIBERNET, is currently investigating a "passive, transmissive" star arrangement, utilizing the ETHERNET protocol. The performance of the fiber optic electronics is given. 32 references

ROBE73 Roberts, L.G., "Dynamic Allocation of Satellite Capacity Through Packet Reservation". National Computer Conference, New York, New York, June 4-8, 1973.

The use of satellite channels for data transmission is examined. A traffic model is derived for comparing four slotted channel access protocols: FDM, TDMA, slotted ALOHA and reservation ALOHA (R-ALOHA). Analytical results for channel utilization, delay and optimum number of reserved slots for R-ALOHA are derived. Finally, the four protocols are compared on a cost/performance basis. 10 references

ROBE78 Roberts, L.G., "The Evolution of Packet Switching". Proceedings of the IEEE, Vol. 66, No. 11, Nov. 1978, pp. 1307-1313.

A historical account of the development of packet switching networks is the basis of this paper. The early pioneers are identified, and the role of standardization in the implementation of public data network services is discussed. The trends of packet satellite, packet radio, and packet voice are projected as the future of packet switching technology. 26 references

SHOC79a Shoch, J.F., Stewart, L., Internetwork Experiments with the Bay Area Packet Radio Network. Xerox Palo Alto Research Center, Palo Alto, California, Report No. SSL-79-4, Feb., 1979, pp. 1-24.

The use of RADIONET (PRNET) as a transit network between two ETHERNETs is developed from initial implementation testing to limited production environment application. The hardware interface and network software functions of the gateways are described. Internetwork throughput experiments conducted for a variety of source and destination configurations revealed a high overhead for processing each packet. In fact, the RADIONET PRU processor was the limiting factor. 7 references

SHOC79b Shoch, J.F., Hupp, J.A., "Performance of an Ethernet Local Network -- A Preliminary Report". *Local Area Communications Network Symposium,* Boston, Massachusetts, May, 1979.

The initial measurement results of an Ethernet local area network are given. Performance is analyzed under three different traffic load conditions: normal, high load and overload. The measurement techniques are discussed, but not presented in detail. 8 references

SHOC79c Shoch, J.F., *An Annotated Bibliography on Local Computer Networks.* Xerox Palo Alto Research Center, Palo Alto, California, Report No. SSL-79-5, Oct. 1979, pp. 1-47.

The bibliography has two parts: the first lists references on local area networks, the second, radio based systems. Each part lists the networks, then the annotated citations follow in alphabetical order by primary author. Scientific and technical papers, as well as product announcements and reports, are included. Not all entries are annotated. 385 references

STOK76 Stokes, A.V., Bates, D.L., Kirstein, P.T., "Monitoring and Access Control of the London Node of ARPANET". *National Computer Conference,* New York, New York, June 7-10, 1976.

The use of the University College London node of ARPANET was monitored to assess the connect time characteristics of host usage. The hardware and software tools are described, and results of the connect time pattern of one user and the dominant host (64% of total usage) are given. Monitoring was not continuous over the time period: significant peak traffic periods were often absent. 9 references

SUNS76 Sunshine, C.A., "Factors in Interprocess Communication Protocol Efficiency for Computer Networks". *National Computer Conference,* New York, New York, June 7-10, 1976.

Protocol parameters that are independent of a network's physical characteristics are examined in terms of their effect on performance statistics. The retransmission interval, flow control strategy, buffering and acknowledgment strategy, and packet size are analyzed with respect to throughput, delay, retransmission rate, line efficiency, and buffer requirements. Mathematical models based on probability and queueing theory are given for each parameter, together with probabilistic data examples. 23 references

SVOB77 Svobodova, L., *Computer Performance Measurement and Evaluation Methods: Analysis and Applications.* New York, New York: American Elsevier Publishing Company, Inc., 1976.

The emphasis of this book is on computer system performance measurement and evaluation methods. It is included in this Bibliography because it is one of the few sources of documentation on the different types of measurement tools. The design concepts, power and limitations, and applications of a computer system monitoring facility, as presented, are extendable to computer network measurement. Other topics covered are: the need for computer system performance evaluation, measures of performance, workload models, simulation and measureability. A measurement case study is also discussed. 315 references, not all unique.

THOR79 Thornton, J.E., "Overview of HYPERchannel". *IEEE Computer Society International Conference: COMPCON Spring 1979,* San Francisco, California, February 26-March 1, 1979, pp. 262-265.

The HYPERchannel, representing the first commercially marketed local computer network architecture is described. The objectives and characteristics of the architecture are given, followed by a brief history and current status of the development company, Network Systems Corporation. 4 references

TOBA76 Tobagi, F.A., Lieberson, S.E., Kleinrock, L., "On Measurement Facilities in Packet Radio Systems". *National Computer Conference,* New York, New York, June 7-10, 1976.

The measurement tools designed for the Packet Radio (PRNET) Network are related to the intended measurement functions. The tools, consisting of cumulative statistics, trace statistics, snapshot statistics, and an artificial traffic generator, measure the performance of the channel access scheme, protocol operation, repeater performance, traffic characteristics, and network global performance. 13 references

TOBA78 Tobagi, F.A., Gerla, M., Peebles, R.W., Manning, E.G., "Modeling and Measurement Techniques in Packet Communication Networks". *Proceedings of the IEEE,* Vol. 66, No. 11, Nov. 1978, pp. 1423-1447.

The major emphasis of the paper is on network modeling techniques. The section on measurements discusses the design of measurement experiments with respect to objectives, goals, and measures for performance evaluation. The role of the measurement facilities, in the form of cumulative statistics, snapshot statistics, trace statistics, artificial traffic generators, data collection and analysis, and network control for reliability monitoring, are put in perspective with the measurement task. Examples from ARPANET, PRNET and SATNET illustrate the concepts. 94 references

TOKO77 Tokoro, M., Tamaru, K., "Acknowledging Ethernet". *IEEE Computer Society International Conference: CompCon Fall 1977,* Washington, D.C., September 6-9, 1977, pp. 320-325.

An improvement is proposed for ETHERNET using, first of all, acknowledgment packets, and secondly, assigning priority to them over user data packets. Simulation results show increased effective throughput and shorter response time when the number of hosts on the ether is large (>50). The algorithm and hardware interface block diagram are also given. 2 references

TOWN77 Townsend, V.S., *Database of References for Computer Performance Modeling, Measurement and Evaluation.* Documentation report prepared for Professor Ursula Maydell in partial fulfilment of the requirements of Computing Science 660, University of Alberta, May, 1977.

The papers presented in a seminar course on computer system performance analysis were compiled in a computerized database using SPIRES. Documentation on SPIRES, the database structure, searching techniques, and file maintenance are contained in the manual. Printouts of the database indexes and main records, plus the file and format definitions are also included. 5 references

TOWN80a Townsend, V.S., Maydell, U.M., *Computer Network Performance Bibliography.* University of Alberta, Department of Computing Science, Technical Report TR 80-2, April, 1980.

The bibliography is a printed version of the NETWORKS online database [see TOWN80b]. At the beginning is a list of all database entries, fully annotated with keywords and complete reference information. An author index, affiliation index, source index and keyword index follow. Preceding each index is a term list with frequency counts.

TOWN80b Townsend, V.S., Maydell, U.M., *User's Manual for the NETWORKS Bibliographic Database on Computer Communication Network Performance Analysis.* University of Alberta, Department of Computing Science, Technical Report TR 80-3, April, 1980.

The User's Manual is the technical documentation for the NETWORKS online database, implemented using the SPIRES database management system at the University of Alberta. Topics covered are database design, online retrieval, file maintenance, SPIRES overview, and the NETWORKS file definition, formats and protocols. The database is also available in printed form [see TOWN80a].

TREA78 Treadwell, S.W., Hinchley, A.J., Bennett, C.J., "A High Level Network Measurement Tool". *European Computing Congress: Eurocomp 78,* London, England, May 9-12, 1978, pp. 35-49.

A tool to assess network performance as seen by an end user is described. The advantages and disadvantages of using real versus artificial traffic sources are compared. The purpose of network versus end to end traffic is contrasted. The software package developed, known as GNOME, uses timestamping of packets as they pass through various layers of protocol and/or points in the end to end route. Bulk, interactive or stream traffic may be generated, and up to 30 network activities can be simulated by each GNOME. The GNOME also handles the statistics collection and experiment control functions. GNOME is implemented in the gateway minicomputers of the SATNET satellite broadcast network. 6 references

ULME78 Ulmer, S.W., "Evaluation and Measurement of the Performance Parameters for a High Speed Computer to Computer Satellite Link". *International Conference on Computer Communication,* Kyoto, Japan, September 26-29, 1978, pp. 623-628.

The emphasis of this paper is on the parameters and measurement subsystem used to determine the performance characteristics of the point to point High Level Data Link Control (HDLC) protocol in a satellite link environment. The performance parameters are classified into four categories: goal oriented, HDLC control, environmental (hardware constrained), and resource consumption. The measurement subsystem consists of a hardware and software monitor, and an offline data reduction program. Results of the experiments using a Symphonie satellite (1.544 Mbps channel) and two IBM computers located in Gaithersburg, Maryland, and La Gaude, France, are not reported. 2 references

USGO79 U.S. Government, "Measurement of Interactive Response Time". *Computer Communication Review,* Vol. 9, No. 1, Jan. 1979, pp. 10-24. [Reprinted from "Federal Information Processing Standards Publication 57".]

Guidelines are given for measuring the quality of interactive computer service delivered to keyboard terminal users through computer networks. The user service measures of response time and interactive turnaround time are defined. The environmental conditions during the measurement process, the driver and monitor configurations, the data analysis and presentation methods are all specified. no references

WEDB74 Wedberg, G.H., Hauschild, L.W., "The General Electric Network Monitor System". *IFIP Congress: Information Processing 74,* Stockholm, Sweden, August 5-10, 1974, pp. 24-28.

The monitor system developed for the General Electric International Network is described. The functional requirements for operational status information, and statistical performance information are given. 2 references

WEIR78 Weir, D.F., Prater, W.E., Dam, X.N., "X.25 Test Facilities on Datapac". *International Conference* on Computer Communication, Kyoto, Japan, September 26-29, 1978, pp. 273-279.

Some tools developed for testing and debugging host interface software to DATAPAC are discussed. For the purpose of verifying an X.25 implementation, an X.25 Protocol Tester was built. Based on a microcomputer system, a user or node can be emulated. Data generation of packets is controlled by parameters. Diagnostic codes were added to the node software to assist in analyzing exception conditions. Finally, a hardware unit for line monitoring is described. 8 references

WEIS78 Weissler, R., Binder, R., Bressler, R., Rettberg, R., Walden, D., "Synchronization and Multiple Access Protocols in the Initial Satellite IMP". IEEE Computer Society International Conference: COMPCON Fall 1978, Washington, D.C., September 5-8, 1978, pp. 356-362.

This paper begins with a description of the physical configuration of the SATNET satellite broadcast network. The emphasis of the paper is on the synchronization procedures used. A global time scale was required for the multiple access protocols being evaluated. A hardware interface, which monitors the propagation time of the messages sent (which varies as the satellite drifts) and the time reference (relative to the other SIMP clocks), ensures that messages sent at the same global time reach the satellite at the same real time. A brief discussion of the slotted channel structure, function of routing packets, and implementation of the acknowledgment scheme follow with an overview of the F-TDMA, R-TDMA, S-ALOHA and R-ALOHA protocols. 17 references

Appendix A -- Overview of SPIRES

In this Appendix, an overview is given of the SPIRES database management system. The material is condensed from the references listed at the end of the Appendix (Section A.7).

The Stanford Public Information Retrieval System (SPIRES), is a generalized, online database management system developed at Stanford University in the early 1970's.

The task of the original SPIRES development was to provide a file management and data retrieval system for Stanford's library automation project (BALLOTS). The generality of the present design can be gauged by the diversity of applications that SPIRES now supports.

SPIRES users design and maintain their own databases; there is no centralized database administrator.

The primary database unit is the subfile: a set of "goal records" optimally linked to "index records".

Records consist of elements, which may be required or optional, fixed or variable in length, singly or multiply occurring. It is possible to group logically related elements into "structures". Structures may contain other structures, nested up to ten levels.

Each record has a unique identifier called the "key". SPIRES can assign record keys if a natural key does not exist as an element.

Indexes for rapid retrieval of records are created and maintained automatically by SPIRES. For unindexed information, sequential search techniques are available. It is also possible to obtain a subset of the file via an index search, and then process the subset sequentially.

Once a set of records has been retrieved, they can be displayed at the terminal or line printer, or put into a file for manipulation by other user programs. Reporting facilities for formatting output are built into SPIRES.

Adding, updating and deleting of records can be done online, or in batch mode. Data can be validated by processing rules. Processing rules can convert numeric data into binary or floating point representation, process dollars and cents data, compress and compact data, translate and codify data, validate lengths and occurrences of data, supply default values, and generate data such as date and time.

New record elements can be added easily to any SPIRES file. Indexes can be erased and rebuilt: the same commands can be used to add new indexes to the file.

A.1 SPIRES Command and Definition Languages

There are various languages that combine to form the unified SPIRES information system for development and use of databases.

- Interactive Command Language:

  The SPIRES interactive command language is used for input and editing of data, record retrieval and record display.

- Protocols Language:

  The protocols language is an extension of the command language allowing packaging of SPIRES/EDITOR/MTS command sequences in a program that can be executed by users.

  This tailoring is particularly important when the end users of a production file have no special training in SPIRES commands.

  Protocols are developed and tested interactively. String manipulation and arithmetic functions as well as condition testing and branching capabilities allow sophisticated interactive dialog and database processing.

- Format Definition Language:

  Any file user can provide formats for input and/or output of any database.

  By means of input formats, a user can be prompted for the record element values, and be given helpful diagnostic messages and reprompts should any error conditions be raised. Also, input formats provide a tool for converting already existing machine readable data to SPIRES-suitable input form. Using commands for arithmetic and string operations, and condition testing and branching, complex algorithms for data validations can be performed.

  Through the use of output formats, reports are produced by mapping file element values, and any other computed values, onto a two-dimensional array that can be output onto a terminal, file, or line printer. Output formats can make SPIRES database contents acceptable for use by batch programs, or can make the same data more easily understood by an untrained user.

- File Definition Language:

  The file definition declares the structure of each database, determines handling of data entered into and retrieved out of the file (input and output processing rules) and specifies indexes to be built by the system.

  The file definition is the most important part of a SPIRES database. It is the realization of the database design, and is the foundation from which the protocols, formats and interactive command language operate.

A.2 SPIRES Program Modules

There are several SPIRES programs and modes.

- Online SPIRES:

  Online SPIRES is the program with which most file searching and updating is done.

- SPICOMP:

  An online program, SPICOMP compiles file definitions, format definitions, variable group definitions and protocols, returning error messages if the compilation process is not sucessful.

- SPIBILD:

  Both a batch and online program, SPIBILD is called to pass records from the goal records to the index records.

- FASTBILD:

  FASTBILD is a program that greatly reduces the CPU time and I/O's necessary for adding a large number of initial records to a file.

- Host Language Interface:

  This facility allows access to SPIRES databases through user programs written in PL/1, COBOL and other languages. The user programs can provide input and/or process output from SPIRES files. (Not currently available at the University of Alberta.)

A.3 Physical Organization

The SPIRES physical records, or blocks, are 2048 bytes in length: this is the amount of contiguous data that can be retrieved from disk into core in one read operation.

SPIRES files are organized into a structure known as a B-tree. Figure 4, redrawn from [JACK77], depicts a very simple B-tree structure. The detailed information maintained in each block can be found in [JACK77].

Each block begins with a header containing block control information which is used to validate the integrity of the data stored in the block. Trailers are located at the

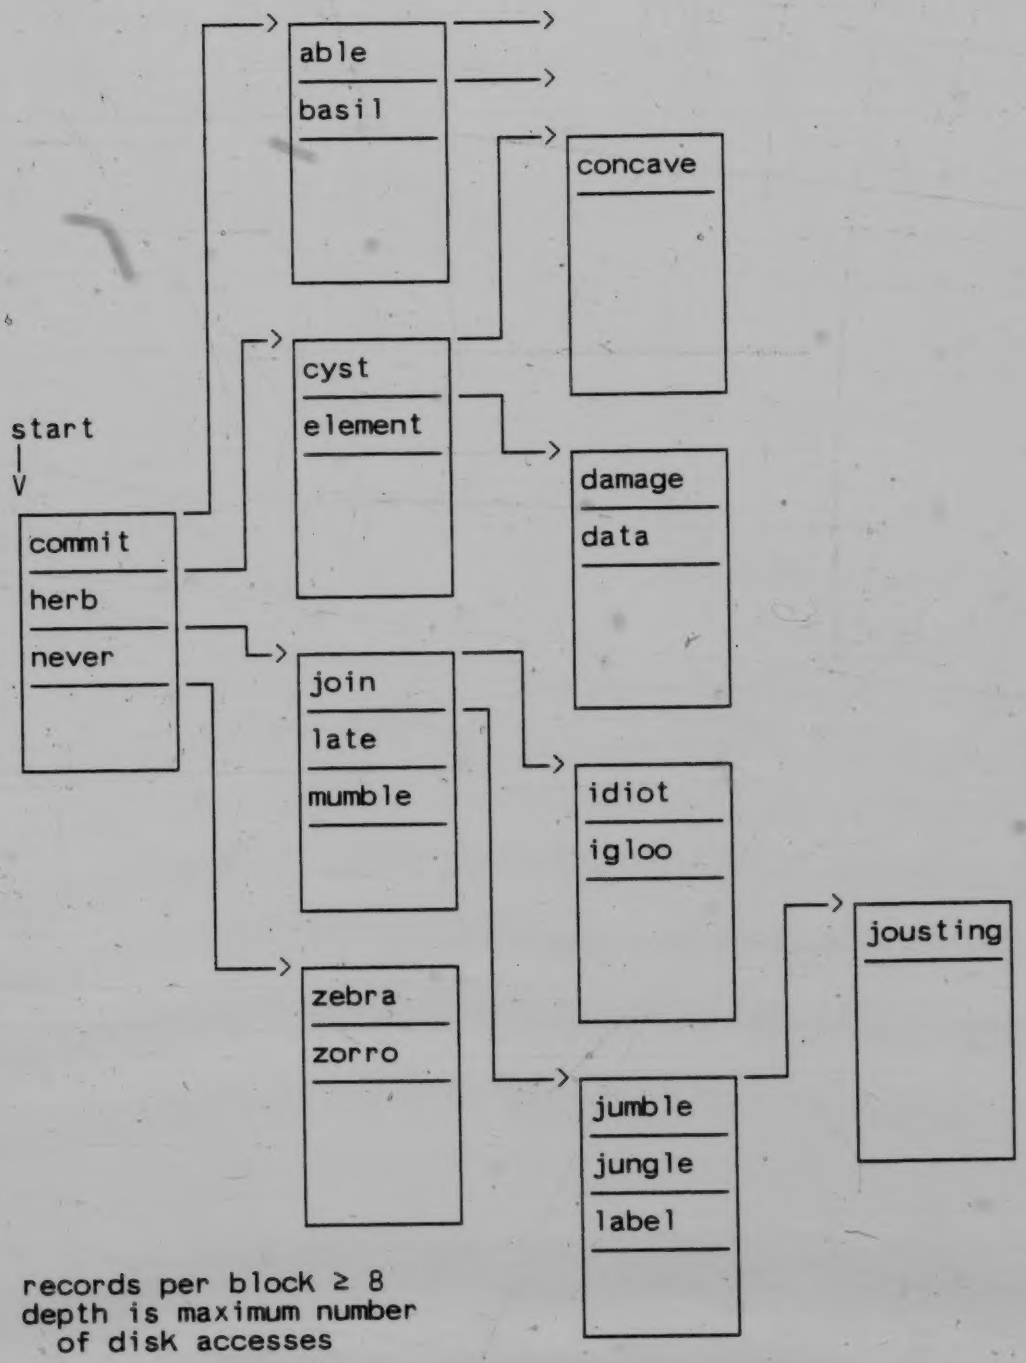
start
able
basil
concave
cyst
element
damage
data
commit
herb
never
join
late
mumble
idiot
igloo
jousting
zebra
zorro
jumble
jungle
label
records per block ≥ 8
depth is maximum number
of disk accesses

Figure 4 - Simplified SPIRES B-tree Structure

end of each block. Trailers are used to give the location, or displacement, of each logical record in the block, and are maintained in key sequence. The actual records are stored after the block control information. Each record is preceded by a branch pointer to a block deeper in the tree.

Retrieval of records begins with the block that is the root of the tree. The block is scanned for the record key value. If found, the search process is complete. Scanning of the block stops with the first key greater than the record in sort sequence. Preceding this key is a pointer to a block in the next level of the tree. The block is read into memory, and the search process is repeated until the record is found.

For access efficiency, SPIRES often "removes" records into a dataset called the residual. The B-tree then contains only the key of the record, and an absolute pointer to its location in the residual. Thus the length of a record does not degrade retrieval efficiency. It is possible for the file definer to request that the removal process be done unconditionally.

During the database update procedure, SPIRES dynamically invokes a localized rebalancing process to maintain a uniform depth for the tree, and to optimize the minimum average accesses per record.

A.4 Privacy and Security

In addition to MTS level account and password validation, SPIRES offers several levels of access security.

Privilege groups are used to define the "rights of access" of a group or individual user. The privilege group determines the user's view of the file and its data, which varies according to the number and variety of restrictions placed on elements and indexes.

Record level access control is provided by "subcodes" which operate on a "security" element. Depending on the bit pattern of the data value, a record can be flagged for access and/or updating privileges.

Data manipulation privileges are controlled by "secure-switches". Secure-switches activate certain functions which may be used independently or in any combination. A few examples are:

- prevent the user from adding, updating or removing records from the database (read only access);
- prevent non-formatted display of records;
- limit retrieval to index and sequential searching; and
- determine the action taken when an error is detected by a processing rule.

Access control at the element level is achieved by assigning "priv-tag" numbers to each element. The file definer can then indicate to SPIRES for each privilege group:

- those elements that cannot be seen;
- those elements that cannot be updated; and
- those elements that cannot be searched.

A.5 Database Recovery

SPIRES uses checkpointing in its file update routine. Also, the module that maintains the deferred queue (containing record updates, additions and deletions) duplicates critical write information to maintain consistency of the data. Automatic recovery procedures are invoked if an error is detected.

In addition to a summarized audit trail of file modifications, SPIRES has the ability to provide "before" and "after" images of all records removed or updated. These can be used to restore a database to its original state in the event of a database disaster.

A.6 File Management Tools

There are several commands available to the file owner for monitoring the status of a database. The type of information provided is:

- detailed information about each record type;
- a summary of the physical state of the file, including blocks used, space utilized, number of records, average number of accesses per record and tree depth;
- activity of the file, including the number of adds, deletes and updates performed and the maximum record length;
- an audit trail of modifications, providing date, time,

user account, terminal type, transaction type and record key; and

- record overhead information.

The file owner can request logging of user accesses and manipulations of a database. Three levels of logging can be specified:

- user accesses to a database;
- all SPIRES commands issued; and
- all protocol prompts and responses.

A database can be defined as a chargeable file, with rates for hourly connect time, record display, and a minimum charge. It is the responsibility of the file owner to process, bill for and collect fees for database use.

Commands are available for processing logged information.

A.7 SPIRES Documentation

JACK77 Jackson, G.R. and Senda, R.E., *SPIRES/370 File Definition,* University of Alberta Computing Services Reference Manual 62, July, 1977. Adapted from the Stanford Center for Information Processing publication *SPIRES File Definition,* by J.R. Sack.

JACK78 Jackson, G.R., *SPIRES/370 Protocol Language,* University of Alberta Computing Services Reference Manual 63, June, 1978. Adapted from the Stanford Center for Information Processing publication *SPIRES Protocol Language.*

SCIP76 *SPIRES/370 Formats Language,* Stanford Center for Information Processing publication, November, 1976.

SEND79 Senda, R.E., *SPIRES Data Base Management,* University of Alberta Computing Services draft document, March, 1979. Prepared in parallel with Stanford Center for Information Processing publication *SPIRES Data Base Management.*

SWAR78 Swartz, F. and Jackson, G.R., *SPIRES Searching and Updating,* University of Alberta Computing Services Reference Manual 29, June, 1978. Adapted from the Stanford Center for Information Processing publication *SPIRES/370 Searching and Updating.*

Appendix B -- Database Design Documentation

The SPIRES database management system (described in Appendix A) is used to maintain an online database, called NETWORKS, of references on computer communication network design and performance analysis. The NETWORKS database, designed and implemented by the author, is documented in this Appendix.

The SPIRES manuals listed in Appendix A should be consulted for standard procedures on database searching, updating and file maintenance.

B.1 File Definition

The file definition specifies the structure of the database and the access privileges of users. There are three parts to a file definition:

- goal record definition;
- index record definition and linkage section; and,
- subfile and access privileges section.

The NETWORKS file definition is now described.

B.1.1 Goal Record Definition

The goal record describes the structure of an entry in the database by defining the attributes of each element in the record.

Elements may be required or optional. Required elements represent data that must be entered for the document record to be accepted by SPIRES. Optional elements are entered only if the context applies. The structures are optional: within a structure, however, some elements are required.

The required elements for the NETWORKS database are:

- Number (alias NU):

  The number is the record key. It is stored as a halfword binary integer. Input processing rules verify that the number is all digits, nonnegative, and within the range of values from 1 to 32,767. On output the number is converted to 6 digits, and zero filled on the left.

- Date-Added (alias DA):

  The date the record was added to the database is automatically supplied by SPIRES, and stored in 4 bytes using the SPIRES internal date representation. On output, the month, day and year are printed: e.g., "May 27, 1980".

- Date-Updated (alias DU):

  SPIRES automatically maintains the last date of record modification. It is stored and printed in the same format as Date-Added.

- Publication-Date (alias PD):

  The date of publication of a document is entered using one of a variety of input formats (American style). The date is verified as being valid: e.g., "February 31" is not accepted. The publication date is stored and printed in the same way as Date-Added.

- Review-Depth (alias RD):

  The review depth indicates the level of detail in which the document was examined. Three levels are used: "brief", "general" and "detailed". The review depth is stored internally as a one byte code, and converted into the text string on output.
- Document-Type (alias DT):

  The type of document is identified as one of: "thesis", "book", "report", "journal" or "conference proceedings". Stored as a one byte code, the document type is converted into the text string on output.
- Identifier (alias ID):

  The identifier consists of the first four (or more) letters of the primary author surname, followed by the last two digits of the publication year. Sufficient alphabetics of the surname are used to uniquely identify the author. In the event of two authors with the same surname, the initials are added. If more than one document is published in the same year, then an "a", "b", etc., is appended to the identifier in corresponding chronological order. On input, the identifier index is searched, and if duplication occurs a warning message is printed.
- Title (alias TI):

  The title of the document exactly as it appears on the manuscript is entered. The leading articles "the", "a" and "an" are typed at the end of the title to allow proper collation in alphabetical title listings. Excess and trailing blanks are automatically trimmed.

The optional elements of the NETWORKS database are:

- Treatment (alias TR):

  The treatment element describes the document from a vocabulary of 20 terms. The area of the performance study is indicated by: "system performance", "system model", "workload", "simulation" and/or "measurement tools". The type of paper is characterized by one or more of: "comparative paper", "theoretical paper", "applied paper", "operational paper" or "background paper". The level of presentation is denoted by: "introductory", "intermediate", "advanced", "overview" and "technical". The type of data analyzed is listed as: "examples", "real data", "simulated data", and "probabilistic data". Finally, "questionable" is used to flag references that contain doubtful approaches or methods. Stored as one byte codes, the treatment words are translated into text strings on output.

- Citations (alias CI):

  The document numbers of the papers cited in the reference list are entered. The numbers are stored as two byte binary integers, and printed in the same format as the Number.

- Keywords (alias KE):

  Free text keywords are assigned to each document to identify the subject areas discussed. Each keyword is checked to ensure that it does not exceed 130 characters in length.

- Abstract (alias AB):

  The abstract of the document is input, together with the number of references cited. Author abstracts are acknowledged using "(Author abstract)" at the end; abstracts modified or written are denoted by "(Annotated abstract)".

- ISSN

  The International Standard Serial Number (ISSN) assigned to a periodical is entered.

- ISBN

  The International Standard Book Number (ISBN) assigned to a monograph is entered.

- Language (alias LA):

  If the document is written in a language other than English, the language element is used.

- Abstract-Number (alias AN):

  If the abstract has a number assigned to it, then it is entered.

- Headings (alias HE):

  For documents from journals or conference proceedings, the headings of the various sections are entered.

- Pub-Country (alias PC):

  The country in which the document was published is entered. This element applies mainly to foreign language publications.

- Journal-Title (alias JT):

  The name of the journal is entered if the document type is a journal article.

- Source-Volume (alias SV):

  The volume number of the publication is entered.

- Source-Number (alias SN):

  The issue number of the publication is entered.

- Source-Pages (alias SP):

  The page numbers of the publication are entered.

- Note (alias NO):

  Special notes about the publication, such as reprints of conference papers, papers based on dissertations, and cross referencing of documents, are entered.

- Author Structure (alias AU):

  Information about each author is entered as a structure. The components of the author structure are:

  - Name (alias NA):

    A required element, the name consists of the author's surname and initials. The name format is validated, and checked to ensure that it does not exceed 130 characters.

  - Affiliation (alias AF):

    The affiliation of the author is optionally entered, together with the city and province/state/country. The entry is checked so that it does not exceed 130 characters.

- Comments (alias CM):

  A free text string, optionally entered, provides additional information about the author. An example is when the research was conducted while on leave of absence.

- Book-Details Structure (alias BD):

  The book-details structure provides publishing information about books. The components of the book-details structure are:

  - Publisher (alias PU):

    A required element, the name of the publisher is entered.

  - Pub-Location (alias PL):

    The city and province/state/country of the publishing company is entered. The pub-location is an optional element.

  - LC-Number (alias LC):

    The Library of Congress Catalog number is optionally entered.

  - Edition-Number (alias ED):

    The edition number, other than the first, of the book is entered.

- Conf-Details Structure (alias CD):

  The conf-details structure identifies a conference. The elements of the conf-details structure are:

- Conference-Name (alias CN):

  A required element, the name of the conference is entered.

- Date (alias DA):

  The dates the conference was held is a required entry. This element usually specifies more than one day: the publication date lists the month and year only.

- Theme (alias TH):

  The conference theme is optionally entered.

- Location (alias LO):

  The city and province/state/country hosting the conference is entered. The location element is optional.

- Sponsoring-Body (alias SB):

  The organization sponsoring the conference is optionally listed.

- Editor (alias CE):

  The editor of the conference proceedings is optionally entered. The editor name is validated for proper format.

- Report-Details Structure (alias RE):

  Detailed information about reports is given. The elements contained in the report-details structure are:

  - Publishing-Org (alias PO):

    The organization name that produced the

report is a required element. The city and province/state/country are also entered.

- Report-Number (alias RN):

  The report number assigned to a publication by the producing organization is entered. The report number is an optional element.

- Thesis-Details Structure (alias TD):

  Information about a thesis is entered as part of the thesis-details structure. The structure components are:

  - Degree (alias DE):

    A required element, the degree lists the type of degree awarded, e.g., "M.Sc.", "Ph.D.", etc.

  - Faculty (alias FA):

    The faculty of the university that awarded the degree is entered. The department name may be listed as well. The faculty is a required element.

  - University (alias UN):

    The name of the university that granted the degree is a required element.

  - Advisor (alias AD):

    An optional element, the supervisor's name is entered.

- Contents Structure (alias CO):

  For reports, theses and books, the table of contents is entered. The components of the contents structure are:

- Contents-Title (alias CT):

  The name of the chapter, part or heading is required.

- Contents-Author (alias CA):

  For documents where the chapters and/or papers are written by different authors, their names may be entered. The names are checked to conform to the standard name format.

- Contents-Pages (alias CP):

  The page number that the chapter or part begins on is an optional entry.

B.1.2 Index Record Definition and Linkage Section

The elements of the goal record are "passed" and organized into inverted indexes to optimize retrieval of data. Two functions are served by index records in the NETWORKS database: online access, and bibliography publication.

To facilitate online retrieval of the references from the database, the following indexes were built. Truncated retrieval using word stems is defined for all indexes.

- Title (alias T):

  The title index is built from the title, contents-title and headings elements. The processing rules first convert special characters to blanks. Each entry is then broken down into individual words, using blanks as delimiters. A stop list is used to eliminate common words, and then an index entry for every remaining word is created.

- Author (alias A):

  The author index is built using the name and contents-author elements. The personal name processing algorithm is used to allow searching by surname, and various initial combinations.

- Source (alias S):

  The university, publishing-org, journal-title, conference-name and publisher elements all contribute to the source index. The same processing rules used for the title index are applied.

- Review-depth, Treatment, Document-type (aliases RD, TR, DT):

  The fixed vocabulary words entered as the review-depth, treatment and document-type elements form this index.

- Date (alias D):

  The publication-date is the searchable element of this index. Range searching is particularly suitable for date scanning.

- Keywords (alias K):

  The keywords element, using the same processing rules as for the title index, are applied.

- Citation (alias C):

  The citations element is the contributing data item of this index. All papers that cite a particular reference can be found.

- Affiliation (alias AF):

  Using the same processing rules applied to the title index, the affiliation element is built into this index.

- Identifier (alias ID):

  The identifier element is built into an index, with the truncated search facility coded.

- Abstract (alias AB):

  An index is built with the abstract element, using the same processing rules as the title index.

In addition to searchable indexes, several records are defined to facilitate publication of the bibliography. These indexes are also searchable: the key difference is that the index entries are composed of element phrases rather than the individual words of each element. Truncated searching using word stems is defined for all indexes.

- Source-Phrase (alias SP):

  The university, publishing-org, journal-title, conference-name, and publisher elements all input to this index.

- Keyword-Phrase (alias KP):

  The keywords element is used to build this index.

- Affiliation-Phr (alias AFP):

  The affiliation element is processed for this index.

B.1.3 Subfile and Access Privileges Section

Two privilege groups have been declared for the NETWORKS database.

The A660 computer account is the file owner, and has all privileges for retrieval and file maintenance.

All other computer accounts have read only access to the database: retrieval of references is allowed, but not record modification.

B.2 Format Definitions

Formats are used to display records in an easy to read manner. Two types of formats were written: online display formats, and publishing formats.

The online display formats are designed to for online printing of search result sets. Records can be printed offline: using the SPIRES "set report" command will cause the printout to be dated, page numbered, and have a header at the top of each page. The amount of record information given with each format varies. All use a 70 character width for output.

- SHORT Format:

  The SHORT format prints the record identifier at the left hand margin. The author names, title and complete reference are given, indented 8 spaces.

- DISPLAY Format:

  The first part of the DISPLAY format is identical to the SHORT format. Next, the fixed vocabulary treatment words assigned to the document are listed after the prefix "Treatment:". Similarly, the free text keywords are given following a "Keywords:" string. Finally, the abstract is printed. The additional information is separated by blank lines and indented 8 spaces to retain the alignment established in the SHORT format.

- FULLDOC Format:

  The FULLDOC format prints the entire record. The name of the record item is left justified and followed by a colon. The data value of the record item is printed indented 16 spaces. A blank line is left between record items.

The publishing formats utilize the full 130 character width provided by line/page printers. Unless otherwise stated, two column output is generated. In report mode, the formats produce listings that are paginated, dated, and with headers at the top of each page to identify the type of printout.

- ABSTRACT Format:

  The ABSTRACT format is identical to the DISPLAY format, but with a column print width of 65 characters.

- TITLE LIST Format:

Used for alphabetical title listings, the TITLE LIST format prints the full page width of 130 columns. The document number is printed at the left margin. Indented 8 spaces is the document title, followed by the author names. On a new line, indented 10 spaces, is the document reference.

- AUTHOR LIST Format:

  The AUTHOR LIST format produces a single spaced list of all authors indexed. Three columns per page are used. Preceding each name is a count of the number of papers penned by the author.

- AUTHORS Format:

  The AUTHORS format prints the Author Index. Following each author name, the SHORT format information for each citation is printed, with references given in reverse chronological order, i.e., most recently published documents first. A blank line separates author entries.

- AFFILIATION LIST Format:

  A single spaced list of all affiliations indexed is printed. Each affiliation is preceded by the number of publications attributed to the organization. Scope notes containing cross referencing information is indented 2 spaces from the affiliation name and begins on a new line.

- AFFILIATIONS Format:

  The AFFILIATIONS format prints the Affiliation-Phr index in the same format as the Author index.

- SOURCE LIST Format:

  The Source-Phrase index is printed in the same format as the Affiliation List by the SOURCE LIST format.

- SOURCES Format:

  The SOURCES format prints the Source-Phrase index in the same format as the Author index.

- KEYWORD LIST Format:

  The Keyword-Phrase index is printed in the same format as the Affiliation List.

- KEYWORDS Format:

  The KEYWORDS format prints the Keyword-Phrase index in the same format as the Author index.

B.3 Protocol Definitions

Protocols allow packaging of MTS/SPIRES/Editor commands in macro like manner. Condition testing and variable assignment enhance the processing capabilities of protocols.

For the NETWORKS database, two protocols were written. Both prompt for the date that is to appear on the listings: the current date is generated as the default. Pagination is continuous, beginning with page 1.

- TERMLISTS Protocol:

  The TERMLISTS protocol produces a printout of all the term lists (author, affiliation, source and keyword).

- NETBIBLIOGRAPHY Protocol:

    The NETBIBLIOGRAPHY protocol produces the printout of the computer communication network design and performance bibligraphy.